

The Social Construction of Ignorance: Experimental Evidence*

Ivan Soraperra[†] Joël van der Weele[‡] Marie Claire Villeval[§] Shaul Shalvi[¶]

December 3, 2021

Abstract

We experimentally study the social transmission of inconvenient information, that is, information about externalities generated by one’s own decision, through the matching of supply and demand for ignorance. We show that willful ignorance arises from two, interacting sources. First, informed senders suppress a substantial amount of “inconvenient” information, driven by the wish to relieve the decision maker as well as the sender’s own preferences for information. Second, about one-third of decision makers avoids senders who transmit inconvenient information (“shooting the messenger”). On aggregate, this leads to assortative matching between senders and decision makers based on preferences for information states. Compared to random matching, assortative matching leads to a stronger expression of individual preferences: selfish senders remain ignorant and donate less, while altruistic senders seek out informative senders and give more. We discuss applications to information sharing in social networks and to organizational design.

Keywords: Social interactions, information avoidance, assortative matching, ethical behavior, experiment

JEL: C91, D82, D83, D91

*The authors are grateful to participants at the Beliefs workshop in Munich, the IMEBESS Conference in Utrecht, the ESA conference in Dijon, and the FELIS CORTEX workshop in Lyon for useful comments. The authors thank Bjorn Bartling for helpful discussions. They thank Q. Thevenet for programming, and J. Benistant, T. Garcia, R. Suchon, and C. Rimbaud for research assistance. This work has benefited from the financial support of the French National Research Agency (FELIS, ANR-14-CE280010-01), from IDEXLYON from Université de Lyon (project INDEPTH) within the Programme Investissements d’Avenir (ANR-16-IDEX-0005) and from the LABEX CORTEX (ANR-11-LABX-0042) of Université de Lyon, within the program Investissements d’Avenir (ANR-11-IDEX-007) operated by the French National Research Agency. It has also been supported by the Van Gogh PHC program. Joel van der Weele thanks financial support from the Dutch Science Association (NWO) via a personal VIDI grant (452-17-004). The research was further supported by a Dutch Science Association grant (NWO VI.Vidi.195.137) and the European Research Council (ERC) grants (ERC-StG-637915 and ERC-CoG-865931).

[†]CREED, University of Amsterdam. Email: i.soraperra@uva.nl.

[‡]CREED, University of Amsterdam. Email: J.J.vanderWeele@uva.nl.

[§]Univ Lyon, CNRS, GATE. IZA, Bonn, Germany. villeval@gate.cnrs.fr.

[¶]CREED, University of Amsterdam. Email: S.Shalvi@uva.nl.

1 Introduction

Many of our decisions impact others, sometimes in a visible way, more often in invisible ways. Such invisible or ambiguous impact occurs notably when our consumption decisions contribute to environmental damage or suffering by humans and animals further down the supply chain. Informing ourselves about these impacts is crucial for making pro-social decisions. However, such information is potentially “inconvenient”, as it may highlight trade-offs between personal profits and moral behavior. A growing literature has shown that people are often reluctant to acquire ethical information, and may actively avoid it to excuse their selfish behavior.¹

Despite growing attention to such willful ignorance, we know little about the interpersonal origins of this phenomenon. In particular, almost all research has focused on the demand for willful ignorance, with much less attention paid to the supply side. However, everyone regularly supplies information to friends, colleagues and contacts on social media, and must sometimes decide whether to truthfully relay inconvenient information. Different motives may play a role in the decision to share such information, such as paternalism and a feeling of duty to tell the truth. Other motives might be more strategic, as people may cater to (perceived) demand for willful ignorance. Thus, supply decisions may depend on whether others reward provision of truthful information, or prefer to seek out those who shield them from inconvenient truths.

In this paper, we investigate these motives for information sharing and their effect on ethical decision making. To do so, we conduct a laboratory experiment in the context of ethical dilemmas with uncertainty. Participants in the role of “decision makers” face a choice between actions that differ in their profitability for the decision maker and carry a risk of entailing a negative externality for a charity. Each decision maker is matched to another subject in the role of “sender”, who is likely to have superior information about the consequences of each decision for the charity. The sender can choose to disclose this information to the decision maker, or send an irrelevant distraction (a picture of a cute animal) instead. In this setting, we study the willingness of senders to supply “inconvenient information”, that is, information that highlights a trade-off between profits and externalities.

¹A number of papers shows that people engage in “willful” or “strategic ignorance” of inconvenient information as an excuse of selfish behavior. The first studies demonstrating this behavior are [Ehrich and Irwin \(2005\)](#) and [Dana et al. \(2007\)](#), followed by fast growing number of replications and follow-ups ([Larson and Capra, 2009](#); [Nyborg, 2011](#); [Conrads and Irlenbusch, 2013](#); [Grossman, 2014](#); [Feiler, 2014](#); [Bartling et al., 2014](#); [Kajackaite, 2015](#); [van der Weele, 2013](#); [Grossman and van der Weele, 2017](#); [Espinosa and Stoop, 2019](#); [Serra-Garcia and Szech, 2019](#)). For an analysis of the neural correlates of information seeking or avoidance, see [Charpentier et al. \(2018\)](#); [Sharot and Sunstein \(2020\)](#). Related work shows how self-serving interpretations of risk and ambiguity increase selfishness in sharing decisions ([Haisley and Weber, 2010](#); [Di Tella et al., 2015](#); [Exley, 2015](#); [Garcia et al., 2020](#)). [Freddi \(2019\)](#) provides evidence of information avoidance from the field. There is also ongoing and inconclusive research about willful ignorance in product markets, with [Bartling et al. \(2015\)](#) finding little evidence, whereas [Momsen and Ohndorf \(2020\)](#) and particularly [Ehrich and Irwin \(2005\)](#) and [Momsen and Ohndorf \(2019\)](#) find more positive evidence.

To study the strategic aspects of ethical information sharing, we vary the possibility for decision makers to *choose* their information sources. Across different treatments, we implement either random matching between informed senders and decision makers, or, like in a social network, give decision makers the possibility to choose senders on the basis of their past information sharing decisions. To mimic senders’ incentives to attract clients or followers on social media, senders earn money from each matched decision maker in each treatment. This treatment allows us to study whether decision makers “shoot the messenger” of bad news, and whether this affects the supply of inconvenient information.

We find that social interactions produce willful ignorance: Despite the salience of the ethical dilemma and the presence of relevant information, a plurality of the content shared in our experiment consists of irrelevant animal pictures. Senders are willing to suppress inconvenient information, with a majority doing this sometimes, and a quarter doing it most of the time. We find evidence that suppression results from paternalistic motives, as senders’ own preferences and attitudes toward information predict what they share with others. We also find a clear demand for ignorance: about one-third of decision makers shifts toward senders who suppress inconvenient information if they have the choice to do so. Despite this, we do not find clear evidence for strategic motives of suppression: while senders’ beliefs about the demand for ignorance are predictive of their supply choices, the role of beliefs is similar in the random matching condition.

This diversity of supply and demand for information and ignorance leads to assortative matching: Compared to random matching, decision makers who can choose their senders are more likely to match with senders with similar preferences for information in the social dilemma, as these are more likely to supply their preferred information content. Thus, decision makers with a preference to remain ignorant act more selfishly as result of assortative matching, while those who like to be informed become more pro-social compared to a setting in which matching is random. As a result of these counteracting dynamics, assortative matching leads to a stronger expression of individual preferences: selfish senders remain ignorant and donate less, while altruistic senders seek out informative senders and give more.

Thus, we show that social interactions facilitate the production of willful ignorance, resulting from a complex set of motives. Our study is explorative and does not test a particular model, but the results speak to various applications where people can choose to share ethical information. The fact that irrelevant distractions make up the plurality of information content is in line with casual observations on social media platforms, and may help explain persistent misperceptions about uncomfortable topics, like climate change, children’s labor, or animal suffering. Our results also provide confirmatory evidence for the segmentation of information into “filter bubbles” (e.g., [Aiello et al., 2012](#)), where users match depending on the type of

information that they like to see. The suppression of ethically relevant information is a novel finding that deserves follow-up investigation, given the ubiquitousness of sharing decisions both on and offline.

A second application relates to organizational design. Effective decision making in organizations requires that advisers and consultants give executive decision makers relevant and unbiased information, rather than being “yes-men”. If we allow some extrapolation, our matching treatments can be viewed as comparing a fixed bureaucracy, or “deep state”, with a system where executives bring in their own advisers. We find that both systems lead to similar amounts of prosocial behavior. However, the matching scheme determines *who* has more power over the decision. Under random matching, the power of the executive is reduced, as we find that the decision makers’ preferences in the social dilemma are less predictive of their decisions in this case.

Our paper makes several contributions to the existing literature. First, we contribute to the literature on individual information avoidance, by studying not just demand, but also the supply of information. Previous studies have considered advice in social dilemmas, such as Schram and Charness (2015) and Coffman and Gotthard Real (2019), where unlike in our study, advisers do not have an informational advantage and can only express their opinion.² Lind et al. (2019) allow subjects in experimental ethical dilemmas to force information on decision makers even if they declined it, and show that this causes more decision makers to inform themselves. Several papers investigate the role of image concerns in the supply of information about externalities (Foerster and Van der Weele, 2018a,b; Bénabou et al., 2018). As far as we know, we are the first to study the matching between the supply and demand for ethical information.³

Second, we relate to a growing literature on group decisions and the dilution of responsibility. While the contributions are numerous, prominent examples include Dana et al. (2007) who document how pairs of subjects are more selfish than individuals, and Falk and Szech (2013) who show that more subjects consent to killing a mouse when there is joint responsibility. Bartling and Fischbacher (2012) show that people can partially avoid responsibility by delegating unkind actions to an intermediary. Weisel and Shalvi (2015) introduce complementarities in unethical behavior in a lying task, and show that lying is more prevalent in teams than in individual decision making. Kocher et al. (2018) find a strong dishonesty shift when individuals decide as group members that is driven by communication within groups. More generally, Charness and

²A number of studies have looked at settings where advisers have an informational advantage, but these studies typically focus on conflicts of interest between advisers and decision makers, looking for instance at the disclosure of such interests (e.g. Ismayilov and Potters, 2013) or cognitive dissonance of advisers (Chen and Gesche, 2016; Gneezy et al., 2020).

³There is a small literature on yes-men that studies the role of incentives in biased transmission of information within organizations (Prendergast, 1993). Opinion conformity with those of a manager has also been identified as a strategy of ingratiation for agents who compete for a promotion (Robin et al. (2014), see also Cummins and Nyman (2013)) Here we consider instead the moral domain and a setting where advisers and decision makers are independent.

Sutter (2012) provide survey evidence that groups make more selfish decisions than individuals. We contribute to this literature by showing how informed and uninformed players collaborate on information suppression and its effect on ethical decision making.

The remainder of this paper is organized as follows. Section 2 introduces our experimental design and procedures and presents our main behavioral conjectures. Section 3 presents and discusses our results. Section 4 concludes.

2 Design, Procedures and Conjectures

2.1 Experimental Design

The experimental design consists of three parts and two treatments. The first two parts are identical for all participants, and serve to elicit some characteristics of interest from each participant and familiarize them with the experimental setting. The third and main part differs across treatments. We describe each part in turn. The instructions are available in Appendix A.⁴

2.1.1 Part 1: Elicitation of Social Preferences and the Demand for Ignorance

The first part is designed to elicit the social preferences of the participants under two successive information conditions. We inform participants that a €15 donation will be made by the experimenter to a charity, GiveDirectly, but depending on their decision, this donation can be cancelled.⁵ Participants have to make a first decision by choosing between two options under complete information. Option 1 pays them €9 and confirms the experimenter’s donation to the charity, while option 2 pays them €15 but cancels the donation, introducing a moral dilemma. Before making their decision, participants can see a picture and a testimonial of a potential recipient of the donation taken from the website of GiveDirectly (see an example of picture in the instructions in Appendix A).

After this decision, participants have to make a second decision that is similar to the first, but under incomplete information, analogous to Dana et al. (2007). This decision gives us a measure of the demand for ignorance in a context in which there is no direct social interaction with others. The program determines randomly whether Option 1 or Option 2 cancels the donation, where either possibility is equally likely. The diagram in Figure 1 provides a summary of the payoffs in the two options. Participants are not informed of the outcome of the random

⁴The experimental design was pre-registered at AsPredicted: <https://aspredicted.org/blind.php?x=dr7743>.

⁵We informed participants that GiveDirectly (<https://www.givedirectly.org>) is a charity that transfers money to very poor families in developing countries and that this charity is rated as one of the 7 top charities in terms of cost-effectiveness by the charity evaluation site GiveWell, above many traditional charities in the world. We also distributed a document on the operating mode of GiveDirectly and displayed information from Wikipedia. We chose this charity because its website allows us to select pictures and testimonials of potential beneficiaries who have passed its screening.

draw. However, before making their choice, they have to choose whether they want to be informed about the consequences of their action for the charity. If they select “Beneficiary”, they learn which option cancels the donation and their screen displays the picture and testimonial of a potential beneficiary before their choice of option. If they select “Cute animal”, they remain uninformed: their screen displays an uninformative picture (a cute animal) and they will never learn the consequences of their action, neither before nor after their choice of option. The display of a cute animal is designed to capture a fun distraction of the kind we often encounter on the Internet, and to balance the use of recipient pictures when subjects receive information about the “Beneficiary”.

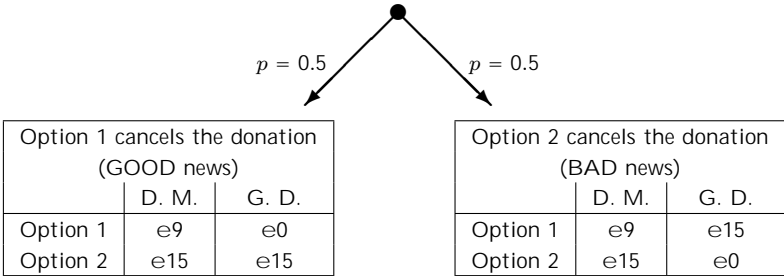


Figure 1: Summary of the decision maker's and charity's payoffs

Notes: The diagram summarizes the payoffs of Option 1 and Option 2 for the decision maker (D.M.) and the charity (G.D.) when Option 1 cancels the donation (GOOD news) and when Option 2 cancels the donation (BAD news). Note that the notions of Good and Bad news were not used in the experiment.

After deciding on being informed or not and before choosing their option, participants are also asked to guess the number of other participants in the session selecting each type of picture. A correct guess pays e1. As we explain below, this elicitation serves to better understand the strategic motives of the senders in the experiment.

2.1.2 Part 2: Role Familiarization

In the second part we familiarize participants with the two roles, senders and decision makers (called “receivers” in the instructions), that will characterize social interactions in Part 3 through two incentivized choices. By introducing only some elements of the more complex environment that will be used in Part 3, this part aims at helping participants to understand the two roles, regardless of the role they will eventually play in the following part. The decision is identical to the decision under uncertainty made in Part 1: Option 1 pays e9 and Option 2 pays e15 to the decision maker, and the program selects randomly for each decision maker which one of the two options cancels the donation to the charity, with a 0.5 probability for each option.

Participants are matched in groups with six other participants. First, all participants play

in the role of a sender. Participants decide whether to send or not information to a decision maker both in the scenario that Option 1 cancels the donation and in the scenario that Option 2 cancels the donation. This gives us information on whether participants are willing to supply information or ignorance to others depending on whether news is “good” or “bad”. Note that throughout the paper, we will refer to news as “good” if Option 1 cancels the donation since in that case, choosing Option 2 maximizes the payoffs of both the decision maker and the charity. We will refer to news as “bad” if the more lucrative Option 2 cancels the donation, since this generates an ethical trade-off between the decision maker’s and the charity’s interests.

After making their choices as a sender, all participants play in the role of a decision maker. Each participant is randomly matched with another player in the group of seven. The decision maker’s information depends on the choice of this other participant when he or she played in the role of a sender. If this sender has decided to share information, the decision maker screen indicates which option cancels the donation and displays the picture and testimonial of a potential recipient before the decision maker’s choice of option. If the sender has decided not to share information or in the case he or she was not informed (with a 0.2 probability)⁶, the screen displays the picture of a cute animal: the decision maker does not know which option cancels the donation and cannot see the picture of a potential beneficiary of the donation. When a decision maker can see the picture of a cute animal on the screen, he or she does not know whether he or she received this picture because the sender selected it or because the sender was uninformed himself or herself. It is common knowledge that if this part is selected for payment, players are paid based on their decision as a decision maker.

2.1.3 Part 3: Endogenous *vs.* Exogenous Matching Between Senders and Decision Makers

In the third and main part of the experiment participants make choices in 25 periods in one of two treatments, RANDOM and CHOICE. We first describe the features that are common to the two treatments. Participants remain matched with the same six other participants as in Part 2. In each group, three participants are randomly assigned the role of sender (identified by a symbol: spade, diamond or club) while the four other participants are assigned the role of decision makers (identified by a letter and a number, R1 to R4). They keep the same role and identifier throughout the part. Decision makers have to choose one of the two same options as in the previous part; Option 1 pays them €9 and Option 2 €15. In each period, the program draws randomly and independently for each decision maker which option cancels the experimenter’s donation to the charity; each option has a 0.5 probability to be selected. Decision

⁶As we explain below, senders are not informed which option cancels the donation with a 0.2 probability and the picture of a cute animal is sent automatically to the decision maker. With a 0.8 probability they are informed. In the second part, we abstract from this aspect, and asked subjects to make choices as if they are informed.

makers are not informed of the outcome of this draw.

By contrast, senders are informed of the consequences of the two options for the charity with a 0.8 probability for each of the four decision makers, independently. The reason that senders are sometimes uninformed is twofold. First, it allows senders some wiggle room, as non-disclosure is not exclusively explained by an unwillingness to share information. Thus, if a sender and decision maker coordinate on the suppression of information between them, they can “pretend” that this ignorance is externally imposed. Second, such wiggle room is a realistic feature of many situations, as sources or advisers are unlikely to always know the truth. In case the sender is informed, the sender has to decide whether sharing or not this information with the decision maker (see screenshots in Appendix A). In the former case, the potential beneficiary’s picture and testimonial are sent to the decision maker with information on which option cancels the donation. In the latter case or when the sender is uninformed, the picture of a cute animal is sent to the decision maker with no information on which option cancels the donation, neither before nor after his or her choice of option.⁷

In both treatments, at the beginning of each period decision makers see a history box that displays a symbol for each type of information sent to him or her by each sender in each of the *previous* periods. Symbols are either ‘GD’ for GiveDirectly—when the sender sent information with the picture and testimonial of a potential recipient—or the symbol of an animal—if the sender had no information or he or she received the information and decided to send the picture of the cute animal (see screenshots in A). The past choices of the sender in the group are only visible to the decision makers, not to the other senders.

The difference between the RANDOM and the CHOICE treatments lies in the matching process of senders and decision makers. In the RANDOM treatment before the decision maker chooses an option, he or she is randomly matched by the program with one of the senders for the current period. He or she receives the information shared by this sender for the current period and chooses one of the two options. In the CHOICE treatment, after observing the history box, each decision maker has to select one of the senders before choosing an option. Thus, in the CHOICE treatment subjects can select which type of sender they prefer, either those who are likely to share information or those who help them remain willingly ignorant.

Senders in both treatments are paid €10 for each decision maker they are matched with, either exogenously in the RANDOM treatment or endogenously in the CHOICE treatment. Senders are not informed about the option eventually chosen by the decision makers, regardless of whether they were matched with them or not. They are informed, however, of the identifier of the decision makers that were matched with them in that period.

⁷Note that the same picture cannot be displayed in more than one period on a participant’s screen.

2.2 Procedures

All sessions were conducted at GATE-Lab, Lyon, France. We ran 16 sessions (8 for the RANDOM treatment and 8 for the CHOICE treatment). The 322 participants (161 in the RANDOM treatment and 161 in the CHOICE treatment) are mainly students recruited from the local engineering, business and medical schools, using Hroot (Bock et al., 2014). 55.3% of the participants are females (57.1% in RANDOM and 53.4% in CHOICE; two-sided Fisher’s exact test, $p = 0.575$). The average age is 22.5 years (22.0 in RANDOM and 23.0 in CHOICE; two-sided Mann-Whitney test, M-W hereafter, $p = 0.131$).⁸ Table B.1 in Appendix B gives a summary of the sessions. The experiment was developed in Java.

Upon arrival, participants drew a tag from an opaque bag assigning them to a computer terminal in the lab. The instructions for each part were distributed and read aloud by the experimenter after completion of the previous part (see Appendix A). Together with the instructions of the first part participants received a description of GiveDirectly and of its operating mode taken from Wikipedia. Before playing the first and third parts, participants had to fill out a comprehension questionnaire. Questions were answered in private. At the end of Part 3 a socio-demographic questionnaire was displayed on the participants’ screen and then they received feedback on their earnings in the session.

The average duration of sessions was 75 minutes. At the end of the session the program randomly selected one of the 28 periods for payment (one of the two decisions in Part 1, the decision as a decision maker in Part 2 or one of the 25 periods in Part 3). If a decision in Part 1 or in Part 2 was selected, participants received either €9 or €15, depending on their chosen option. If a period in Part 3 was selected, the decision maker earned either €9 or €15, depending on the chosen option in that period; the sender earned €10 for each decision maker he or she was matched with in that period (thus, the sender minimally earned €0 if he or she was not matched to any decision maker in that period, and maximally earned €40 if he or she was matched with all four decision makers). GiveDirectly received a donation of €15 for each decision maker whose decision did not cancel the donation. The average payoff of the participants was €18.49 (standard deviation, S.D. hereafter, = 6.57), including a €5 show-up fee. Payments were made in cash, in a separate room and in private.

2.3 Behavioral Conjectures

We designed our experiment to study the drivers of supply and demand for ethically relevant information and their matching in a market, as we explain in more detail below. We consider this an initial exploration, as it is hard to capture our setting in a formal model for two reasons. First, there is a multitude of psychological motives involved in both demand and supply of

⁸Except if specified otherwise, all the non-parametric tests reported in the paper are two-sided and take each individual as one unit of independent observation.

ignorance, relating to self-interest, social preferences, self-image, paternalism etc. Second, the dynamic nature of our interactions allows for a multitude of strategies. Nevertheless, we form a number of conjectures about the behavior we expect to observe, informed by recent literature on willful ignorance.

The behavior of subjects in the experiment will depend on their preferences and their attitudes toward information. Thus, we distinguish between different types of subjects depending on their motivations. First, subjects may differ in their preferences over the payoffs to themselves and the charity. *Selfish* subjects are only motivated by the maximization of their individual payoff, whereas *Altruistic* decision makers are willing to make sacrifices on behalf of the charity. A measure of these preferences is provided by the first decision in Part 1 of the experiment, where subjects make individual decisions with full information.

Second, there are further psychological motivations like self-image and guilt that may affect information attitudes. Previous literature has shown that roughly one third of subjects can be classified as “reluctant” altruists (Dana et al., 2006, 2007; Lazear et al., 2012). These agents want to choose the selfish option, but also want to maintain a positive self-image or avoid cognitive dissonance from being explicitly selfish. Remaining uninformed may serve as an excuse and help maintain self-image while also reaping a profit (Grossman and van der Weele, 2017). The second decision in Part 1 of the experiment shows whether people prefer to avoid information, and allows us to classify subjects as information *Avoiders* or *Seekers*.

How do such heterogeneous motives affect decisions in the experiment? We first consider information supply. Altruistic senders who care about the charity would always want to disclose bad news in order to inform the decision maker of the possible trade-off. By contrast, Selfish senders may deviate from full disclosure for several reasons. Given that senders are paid for each matched decision maker, there may be a strategic motive. If senders anticipate a sufficient demand for ignorance in the CHOICE treatment, they may suppress bad news in order to attract more decision makers. Thus, we expect that in the CHOICE treatment, senders’ suppression will be correlated with their beliefs about the decision makers’ preferences for information. This should not occur in the RANDOM treatment, where matching is random and there is no competition among senders.

Conjecture 1. *(Supply side). There is a strategic supply of ignorance. Selfish senders who believe there is demand for ignorance are more likely to suppress news in the CHOICE treatment.*

There are additional motives for suppression that lead to further hypotheses. For instance, senders may want to do the decision makers a “favor” by not confronting them with a difficult ethical decision. In this case, we would expect to see a correlation with beliefs about ignorance

demand also in the RANDOM treatment. Furthermore, if the sender wants to impose his or her own information preference on the decision maker (paternalism), one would expect that Avoiders are more likely to suppress news. Finally, in all cases, senders should suppress *both types* of news (good and bad ones) in order to really shield decision makers from information. Consistently suppressing only bad news will lead a Bayesian decision maker to infer that uninformative pictures signal bad news.

We now turn to information demand by decision makers. Selfish decision makers should choose Option 2 in all treatments. Because information does not change their decision, standard economic theory would predict that Selfish participants are indifferent between information sources. However, Selfish Avoiders, following the guilt-avoidance logic outlined above, are motivated to select uninformative senders that sent animal pictures in previous periods. By contrast, Altruistic decision makers should choose Option 2 if and only if they are informed that Option 1 cancels the donation, in order to avoid the risk of cancelling the donation by their decision. Indeed, having information about the consequences is necessary to be altruistic, since in the absence of information, either option is equally likely to cancel the donation. Thus, in the CHOICE treatment Altruistic decision makers should select informative senders, *i.e.*, those who in the past periods were most likely to disclose information about the charity. By contrast, Altruistic Avoiders, who would like an excuse in order to behave selfishly, should select uninformative senders. This analysis leads to the following behavioral conjecture.

Conjecture 2. (*Demand side*). *There is a demand for ignorance. Decision makers who avoided information in Part 1 (“Avoiders”) will seek out senders with uninformative messages.*

We can also make some conjectures about differences in matching patterns across treatments. First, if Avoider type senders are more likely not to disclose information, and Avoider type decision makers are more likely not to look for information, we should see assortative matching of types, even if preferences are not directly observable to the other side.⁹

Conjecture 3. (*Assortative matching*). *There is assortative matching in the CHOICE treatment, with Avoider decision makers more likely to match with Avoider senders.*

Finally, Dana et al. (2007) show that if people are able to avoid information, unethical behavior increases. Thus, we expect that unethical behavior will increase in the CHOICE treatment, where senders have an additional reason to suppress inconvenient information (Conjecture 1), and decision makers can select non-disclosing senders.

⁹This conjecture was not included in the pre-registration, and was conceived after the data were gathered. The analysis should therefore be considered more exploratory compared to the other conjectures.

Conjecture 4. (*Ethical behavior*). *In situations with an ethical trade-off, there is more selfish behavior by decision makers in the CHOICE than in the RANDOM treatment.*

3 Results

We first give an overview of the type of information transmitted in both treatments and its impact on decisions. We then turn to analyze the supply and demand of information in the two treatments, as well as the matching of different types of decision makers and senders in the CHOICE treatment and the impact on their behavior. Throughout, we will use “good news” to refer to messages that show no ethical trade-off (Option 1 cancels the donation), “bad news” for messages that show such a trade-off (Option 2 cancels the donation), and “no news” to uninformative animal pictures.

3.1 Overview of Information Consumption and Ethical Behavior

The expected distribution of information available to senders is 40% good news, 40% bad news and 20% no news.¹⁰ Thus, if senders transmitted all information or if decision makers selected only senders who did so, this should be the distribution of information consumed by the decision makers in Part 3. The left panel of Figure 2 shows that the actual distribution of information observed by the decision makers differs starkly from this benchmark ($\chi^2(2) = 594.16$ and $\chi^2(2) = 668.54$ in the CHOICE and RANDOM treatments, respectively).¹¹ With a prevalence of 40.3% in the CHOICE treatment and 41.6% in the RANDOM treatment, the no news constitutes the modal condition of news consumed in both treatments. By contrast, both good and bad news are observed by decision makers about 30% of the time. As we discuss in more detail below, the distribution of consumed information is similar across treatments.

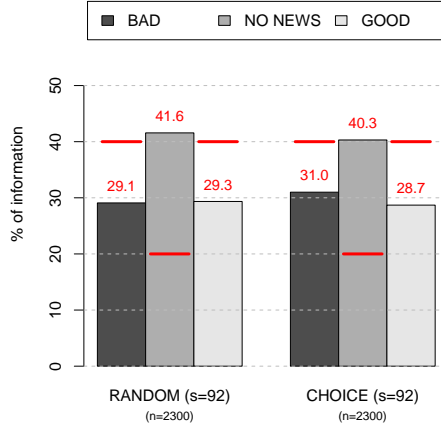
Does it matter what information decision makers consume? The right panel of Figure 2 displays the fraction of choices for Option 2 in Part 3, by information condition and treatment, and shows that it matters a lot. In both treatments, decision makers systematically choose Option 2 after no news (95.5% in CHOICE and 92.2% in RANDOM) or good news (98.2% in CHOICE and 96.3% in RANDOM). Since there is no explicit ethical trade-off in these cases, this shows that subjects understand the choices in front of them. By contrast, when decision makers get bad news, only about 40% of their choices are selfish (39.7% in CHOICE and 40.1% in RANDOM).

Finally, as an overall measure of ethical performance, we consider decisions in the “bad” state, that is, where there was an ethical trade-off (regardless of whether decision makers were informed or uninformed). Using the average amount of selfish decisions per individual as one observation, we find very similar fractions of 62.4% and 62.2% of selfish decisions in the CHOICE

¹⁰The realized frequencies are: 38.3%, 40.7%, and 21.0% in the CHOICE treatment; and 39.7%, 40.8%, and 19.5% in the RANDOM treatment.

¹¹Repeating the test using the realized frequencies instead of the theoretical ones gives the same results.

A - Distribution of information consumed



B - Choice of the selfish option




Figure 2: Information consumption and choices of the decision makers

Notes: Panel A displays the distribution of information observed by the decision makers in Part 3, split by treatment. The horizontal lines show the distribution of information available to senders. Panel B displays the fraction of times Option 2 has been chosen by decision makers, split by treatment and information received. Vertical bars are standard errors based on a linear probability model with errors clustered at group level. Labels below the bars indicate both the number of subjects (s) and the total number of choices (n).

and RANDOM treatment respectively (Kolmogorov-Smirnov test, $p = 0.771$). Thus, we do not find evidence supporting Conjecture 4: the treatment does not affect the aggregate amount of pro-social or selfish behavior by the latter. However, as we show in Section 3.4, this result hides the fact that different groups of decision makers react in opposite ways to the possibility of choosing their own sender.

3.2 Supply of Information

To analyze the supply side, we focus on the suppression of bad news as this is the only news with ethical relevance. To measure the intensity of bad news suppression, we define a sender specific statistic: the fraction of bad states observed by the sender that were not transmitted to the decision maker in the 25 periods of Part 3. Using this metric, we find that only 1% of senders suppresses all bad news, while 29% of senders suppress no bad news at all. Thus, the large majority of senders suppresses bad news at least sometimes, while about 25% of senders suppress more than half of the bad news they receive.

If we use this metric to compare the distributions of sender suppression, we find no statistical differences between treatments (Kolmogorov-Smirnov test, $p = 0.248$). Thus, contrary to Conjecture 1, we do not see an increase in the suppression of bad news in the CHOICE treatment, where senders compete for decision makers and may try to satisfy any demand for information avoidance. Figure C.1 in Appendix C gives an overview of the distribution of sup-

ply choices across treatments. Appendix C also shows individual examples to illustrate various patterns of information suppression by senders with different suppression strategies.

Motives for suppression of bad news. As discussed in our conjectures, various motives may drive the suppression of bad news. To understand these motives, we first consider the role of beliefs about the demand for ignorance. In Part 1 of the experiment subjects reported their belief about the number of other participants in the session that preferred not to disclose information. If senders aim to attract more decision makers, these beliefs should inform their disclosure strategies in the CHOICE treatment. Furthermore, if senders want to do decision makers a “favor”, then we should also see a correlation between beliefs and suppression in the RANDOM treatment. Table 1 shows regression evidence testing these motives. In column 1, we regress suppression of information on sender’s beliefs, and find a highly significant correlation. In column 2 we introduce a treatment dummy, and in column 3 an interaction of beliefs with the treatment. The interaction term is negative and insignificant, showing that if anything, beliefs play a smaller role in the CHOICE treatment.¹²

These results are consistent with the idea that senders are trying to do decision makers a “favor”, rather than strategically adjust their behavior to increase their chance of being selected by decision makers. However, beliefs may also reflect a “false consensus effect” and be a proxy for senders’ own preferences for information. If so, our findings could indicate “paternalism”: a wish to impose the sender’s preferred information or decision on the decision maker.¹³ To see if this projection of preferences explains the effect of beliefs, we control for the sender’s preferences in column 4 of Table 1. Although the confidence level and the size of the coefficient decline somewhat, the coefficient on the belief variable remains significant, indicating at most a modest role of projection. This gives some support to the “doing a favor” type of explanation. In addition, both the coefficients for Selfish and Avider preferences are significant and sizable, increasing suppression by 13.9 and 15.7 percentage points, respectively. This last result indicates a role for paternalism, which is not moderated by beliefs about the decision maker’s preferences.

¹²There are also a number of possible explanations for the lack of a belief effect in the CHOICE treatment. First, it could be due to a different distribution of beliefs across treatments. However, a Kolmogorov-Smirnov test cannot reject the hypothesis of equality of the distributions of beliefs ($p = 0.600$), so this seems unlikely. Second, decision makers could have updated their beliefs during the group phase, on the basis of their experiences. However, repeating the regression analysis presented in Table 1 using only data of the first 5 or 10 periods, yields very similar results. A third explanation is that senders lack the information to be strategic. To reduce complexity, we did not tell senders about the choices of the decision makers nor about the behavior of the other senders. This made it more difficult for senders to optimize their strategy. To speed up learning about demand, senders may have experimented with different strategies, reducing the correlation between beliefs and suppression. We cannot test this explanation within our data-set, but it could be addressed in future research.

¹³In particular, recent results show that many people engage in *ideals-projective paternalism* (Ambuehl et al., 2019), i.e., they restrict others’ choices according to their own preferences. Bartling et al. (2020) show that Americans are willing to intervene in the choices of others when it comes to providing information, but less so when it comes to their choices. Our setting providing an intermediate case, as here information is necessary to make an informed decision.

Table 1: Suppression of bad news by senders

	Model 1	Model 2	Model 3	Model 4
	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)
(Intercept)	0.115 (0.059) [◦]	0.115 (0.065) [◦]	0.079 (0.071)	0.086 (0.072)
Belief # ignorant	0.024 (0.006) ^{***}	0.024 (0.006) ^{***}	0.031 (0.009) ^{***}	0.024 (0.009) [*]
d(CHoice)	---	-0.000 (0.052)	0.070 (0.070)	0.051 (0.069)
d(CHoice) Belief # ignorant	---	---	-0.015 (0.010)	-0.019 (0.010) [◦]
Selfish	---	---	---	0.139 (0.057) [*]
Avoider	---	---	---	0.157 (0.064) [*]
<i>Age</i> <i>Age</i>	-0.003 (0.003)	-0.003 (0.003)	-0.003 (0.003)	-0.001 (0.003)
<i>d(Male)</i>	0.003 (0.046)	0.003 (0.045)	-0.004 (0.043)	-0.043 (0.042)
<i>BAC</i> <i>BAC</i>	-0.009 (0.011)	-0.009 (0.011)	-0.010 (0.011)	-0.013 (0.010)
<i>Number of past participations in exp.</i>	-0.011 (0.017)	-0.011 (0.017)	-0.009 (0.017)	-0.006 (0.017)
<i>Period dummies</i>	YES	YES	YES	YES
Number of observations	5389	5389	5389	5389
Number of clusters	46	46	46	46

Notes: These regressions are based on linear probability models. The binary dependent variable is the sender’s choice to suppress bad news in Part 3. Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. “Belief # ignorant” is the subject’s belief about the number of participants in their session that were willing to remain uninformed in Part 1. Period dummies are included with period 1 as the reference category. *** 0.001; ** 0.01; * 0.05; ◦ 0.1.

To gather further evidence for senders’ motives, we complement the behavioral data by looking at the closing questionnaire, where senders answered the question “According to which principle(s) did you decide to report or not the consequences to the receivers?” As we detail in Appendix F, the answers provide evidence for all the motives discussed above. As examples of doing favors, some senders justify suppression by mentioning that they aim to “relieve the conscience” of decision makers or “make their decisions easy”, and that they tried to “anticipate their expectations”. As examples of paternalistic behavior, some suppressed information depending on whether it was “more profitable for them [the decision makers] to know it or not.” Senders who sent information frequently cite the importance of giving decision makers a choice to donate and exercise their autonomy, while a few mention the wish to attract more clients.

Finally, although the suppression of good news does not affect the decision makers’ choice, senders may suppress information symmetrically to avoid that decision makers infer that “no news means bad news”. In Appendix C, we document that symmetric suppression is indeed common, suggesting that senders think about the inferences that senders are going to make.

In summary, our analysis provide only partial support for Conjecture 1.

Result 1: *The majority of senders sometimes suppress bad news. There is no evidence that competitive motives increase suppression in the CHOICE treatment. Suppression arises at least partially to do a favor to the decision maker and as a result of paternalistic motives.*

Table 2: Suppression of bad news by senders

	Model 1	Model 2	Model 3
	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)
(Intercept)	0.237 (0.062)***	0.174 (0.070)*	0.078 (0.076)
Selfish DM	0.019 (0.038)	0.017 (0.031)	0.008 (0.035)
Avoider DM	0.010 (0.040)	0.013 (0.039)	0.016 (0.040)
Selfish Sender	—	0.229 (0.099)*	0.199 (0.097)*
Avoider Sender	—	0.219 (0.103)*	0.133 (0.100)
Belief # ignorant	—	—	0.023 (0.010)*
d(CHOICE)	0.023 (0.058)	0.012 (0.060)	0.070 (0.066)
d(CHOICE) Selfish DM	0.039 (0.048)	0.038 (0.043)	0.029 (0.046)
d(CHOICE) Avoider DM	0.071 (0.054)	0.048 (0.053)	0.076 (0.053)
d(CHOICE) Selfish Sender	—	0.133 (0.113)	0.106 (0.112)
d(CHOICE) Avoider Sender	—	0.017 (0.133)	0.041 (0.137)
d(CHOICE) Belief # ignorant	—	—	0.019 (0.012)
<i>Age</i> \overline{Age}	0.002 (0.004)	0.001 (0.002)	0.001 (0.003)
<i>d(Male)</i>	0.012 (0.048)	0.045 (0.045)	0.044 (0.041)
<i>BAC</i> \overline{BAC}	0.009 (0.010)	0.012 (0.009)	0.011 (0.010)
<i>Number of past participations in exp.</i>	0.008 (0.016)	0.007 (0.016)	0.006 (0.016)
<i>Period dummies</i>	YES	YES	YES
Number of observations	5389	5389	5389
Number of clusters	46	46	46

Notes: These regressions are based on linear probability models. The binary dependent variable is the sender’s choice to suppress bad news in Part 3. Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. “Belief # ignorant” is the subject’s belief about the number of participants in their session that were willing to remain uninformed in Part 1. Period dummies are included with period 1 as the reference category. *** 0.001; ** 0.01; * 0.05; ° 0.1.

Table 3: Suppression of bad news by senders

	Model 1	Model 2	Model 3
	CHOICE	RANDOM	ALL DATA
	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)
(Intercept)	0.208 (0.079)**	0.274 (0.078)***	0.237 (0.062)***
Selfish DM	0.020 (0.031)	0.016 (0.034)	0.019 (0.038)
Avoider DM	0.086 (0.040)*	0.015 (0.038)	0.010 (0.040)
d(CHOICE)	—	—	0.023 (0.058)
d(CHOICE) Selfish DM	—	—	0.039 (0.048)
d(CHOICE) Avoider DM	—	—	0.071 (0.054)
<i>Age</i> \overline{Age}	0.0001 (0.003)	0.024 (0.014)°	0.002 (0.004)
<i>d(Male)</i>	0.041 (0.055)	0.014 (0.073)	0.012 (0.048)
<i>BAC</i> \overline{BAC}	0.0004 (0.013)	0.014 (0.018)	0.009 (0.010)
<i>Number of past participations in exp.</i>	0.0004 (0.021)	0.011 (0.026)	0.008 (0.016)
<i>Period dummies</i>	YES	YES	YES
Number of observations	2735	2654	5389
Number of clusters	23	23	46

Notes: These regressions are based on linear probability models. The binary dependent variable is the sender’s choice to suppress bad news in Part 3. Robust standard errors clustered at group level are in parentheses. *d* for dummy variables. Period dummies are included with period 1 as the reference category. *** 0.001; ** 0.01; * 0.05; ° 0.1.

3.3 Demand for Information

We now turn to the demand for information in the CHOICE treatment, where decision makers could choose a sender. To test Conjecture 2, we examine the likelihood to choose different senders depending on their profile of past information supply, which is available to the decision maker. To summarize the sender’s information profile, we rank senders according to the relative level of ignorance they provided to the decision maker in the previous 10 periods, i.e., the ranking in period t is based on the number of times senders disclosed information in periods $t-1$ to $t-10$. Then, we ask how frequently the decision makers chose the sender providing the highest, the intermediate, and the lowest level of ignorance. Note that this approach excludes the first 10 periods from the analysis, as senders have not yet established a history.

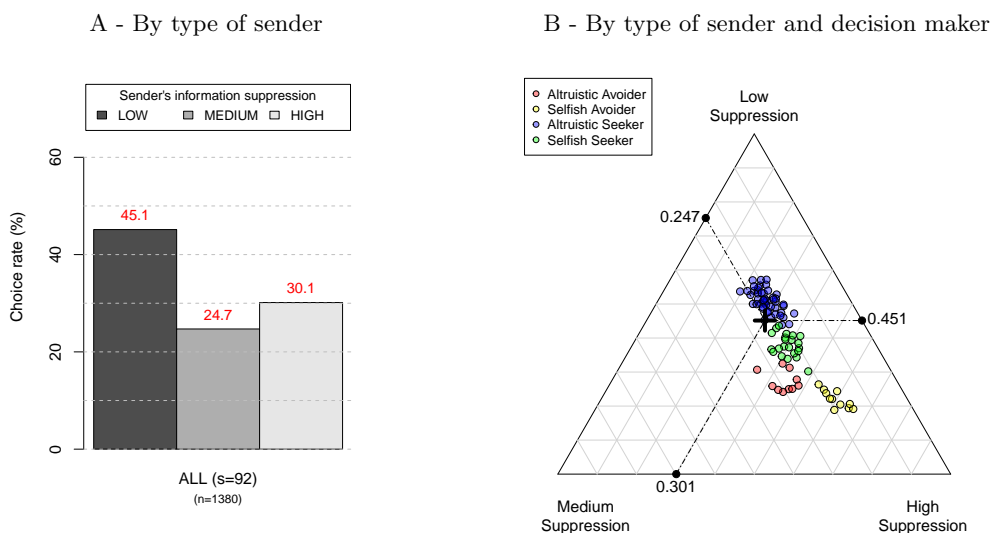


Figure 3: Predicted probability to choose a given sender

Notes: The figure displays the frequency of choices of the three senders in the CHOICE treatment. Senders are ranked (low, medium, and high) according to the relative level of ignorance they provided in the previous 10 periods. Panel A shows displays the predicted probability to choose the sender providing the highest, intermediate, and lowest level of suppression for each of the 92 decision makers. Predicted probabilities are based on Model 1 of Table D.1 of appendix D. Panel B adds the type of the decision maker. The color of the dot captures the type of the decision maker based on decisions in Part 1. The cross shows the average frequency of choice.

Figure 3 shows the frequencies with which different types of senders are chosen. Panel A shows that, on aggregate, the modal choice is the sender that provides the highest level of information. Panel B shows a simplex with the predicted probability to choose each sender for each decision maker based on a multinomial logit model—Model 1 in Table D.1 of Appendix D—where the three alternatives are the senders providing Low, Medium, and High ignorance and the individual specific explanatory variables include the dummies capturing the type of the

decision maker, *i.e.*, Selfish - Altruistic and Avoider - Seeker, obtained from the choices in Part 1.

Predicted probabilities reveal that the aggregate results hide a lot of heterogeneity. The largest group (Altruistic Seekers) clearly penalizes non-informative senders: the most informative sender in the group is chosen, on average, 54.6% of the times whereas the least informative sender is chosen, on average, 22.6% of the times. By contrast, a smaller group (Selfish Avoiders) does the opposite: they choose the most informative sender 27.3% of the times and the least informative sender 44.8% of the times. The simplex shows that the effect of heterogeneity is mostly captured by shifting the probability mass from the Low ignorance to the High ignorance sender, while the predicted probability to choose the Medium ignorance sender is about 20-25% and does not change much across types. Statistical support for these results is reported in Table D.1 of Appendix D.

Thus, our results show that subjects who avoided information in Part 1 are likely to seek out uninformative senders. Finally, in Appendix D, we provide an additional set of measures of information demand. Most importantly, we quantify the degree of information seeking or avoiding of each individual decision maker by comparing their consumption of bad news to the average amount of such news that is available from senders. As we show in the Appendix, roughly 40% of decision makers consume less bad news than the available average in the market, and can thus be classified as Avoiders. The appendix also provides illustrations of different patterns of selection of senders by individual decision makers.

Overall, this analysis supports Conjecture 2 and is summarized in the following result:

Result 2: *A majority of decision makers searches for informative senders. However, a sizable minority of decision makers does not. In particular, subjects who avoid information in Part 1 seek out non-informative senders in Part 3.*

3.4 Assortative Matching and Ethical Behavior

To evaluate assortative matching, we use the two dimensional type classification explained above, based on Part 1 choices. In particular, we look whether decision makers who avoid or seek information in Part 1 are more likely to match with a sender with similar preferences. Panel A in Figure 4 shows the frequency of choice of senders who avoided disclosing information when deciding in Part 1, by decision maker's type. Frequencies are calculated using data of groups where decision makers have the opportunity to choose either type of sender, *i.e.*, groups where there was at least one Avoider and one Seeker among the senders.

The figure shows that (i) being an Avoider substantially increases the probability of matching with an Avoider-sender and (ii) this effect is stronger for selfish decision makers. Indeed, the probability to choose an Avoider-sender increases by 14 percentage points (from 34.6% to 48.6%) for Altruistic decision makers and by 23.3 percentage points (from 26.0% to 49.3%) for

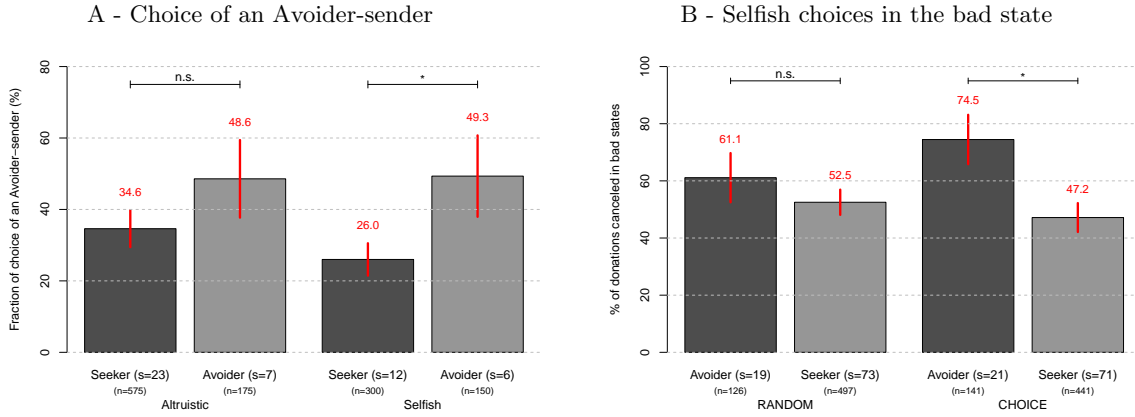


Figure 4: Assortative matching and Selfish decisions in the bad state by condition

Notes: Panel A displays the frequency of choice of an Avoider sender. Decision makers are split by type, as defined based on decisions in Part 1. The frequencies are calculated on the subset of groups where there is at least one and at most two Avoiders among the senders. Panel B displays the fraction of selfish decisions in the different experimental conditions. Decision makers are split into Avoiders and Seekers of information, based on their decision in Part 1 under uncertainty. The fraction of selfish decisions in the CHOICE and RANDOM treatments are computed using cases where the state is bad and all senders received information. In both panels, vertical lines represent standard errors based on a linear probability model with clustering at the group level. In both models, the two factors on the x-axis and their interaction are the only explanatory variable. Pairwise comparisons reported above the bars are based on a Wald test performed using these estimated models. Signific. codes: * $p < 0.05$; n.s. means $p > 0.05$.

Selfish decision makers. This pattern, which receives statistical support from the regression results reported in Model 1 and 2 of Table E.1 in Appendix E, is coherent with the fact that being an Avoider predicts both the demand for ignorance as the supply. The analysis leads to our third result that supports Conjecture 3.

Result 3: *Endogenous matching leads Avoider-type decision makers to match significantly more with Avoider-type senders.*

How does assortative matching affect the behavior of decision makers? In Section 3.1, we already noted that there is no aggregate change in the prosocial behavior of the decision makers between treatments, contradicting Conjecture 4. However, this null results hides some interesting variation, as assortative matching has differential effects on different types of decision makers. In particular, one would expect that the Avoider type become less prosocial in the CHOICE treatment, as they are now less likely to obtain information, while the opposite holds for the Seeker type.

This conjecture is indeed borne out in the data. Panel B in Figure 4 shows the fraction of decisions that cancelled the donation. Here, we condition on being in the bad state and we focus only on the cases where all senders were informed to eliminate those cases where ignorance is

outside the control of our subjects.¹⁴

The first two bars reflect decisions in the RANDOM treatment, split by Avoiders and Seekers. Selfish decisions are 8.6 percentage points higher among Avoiders than Seekers. The next two bars show the corresponding decisions in the CHOICE treatment. We observe that both groups move in opposite directions, with Seekers becoming about 5 percentage points more prosocial, while Avoiders become 13.4 percentage points more selfish, widening the gap between the two groups to 27.3 percentage points. Statistical support for the effects suggested by the figure are provided by the regressions in Model 3 and 4 of Table E.1, Appendix E. Interestingly, the behavior of decision makers in the CHOICE treatment is very similar to their behavior in Part 1 of the experiment. This further underscores that having the CHOICE of senders allows decision makers to remain in control of their decisions.¹⁵

This analysis leads to our last result.

Result 4: *There is no evidence for differences in ethical behavior across treatments on aggregate. However, assortative matching in the CHOICE treatment leads Avoiders to become more selfish than Seekers compared to when matching is random.*

The last two results show that when subjects can choose their senders, their behavioral patterns reflect their own preferences for information as well. In particular, the heterogeneity of sender behavior means that decision makers can obtain the information or non-information they prefer. This in turn impacts their behavior, and behavioral differences between types become pronounced only in the CHOICE treatment.

4 Conclusion

We have shown how social interactions can produce willful ignorance through the behavior of both sides of the interaction. First, senders are willing to share irrelevant distractions instead of relevant information, a decision driven both by their own preferences and their beliefs about the sender’s preferences. Second, about one third of decision makers prefer ignorance, and actively seek out uninformative senders. We find that the resulting assortative matching allows subjects who prefer ignorance to act more selfishly, and those who prefer information to become more

¹⁴When senders are not informed, harm to the charity cannot be avoided or reduced by the actions of the participants. Also, by focusing on situations where information was available, we can compare the decisions with individual decisions in the first part of the experiment. When including observations where senders were not informed, results show the same pattern but are not statistically significant as the data become more noisy.

¹⁵In the individual decisions under uncertainty in which the bad state was realized, Seekers behave selfishly in 45.2% of the cases compared to 47.2% in the CHOICE treatment, while Avoiders behave selfishly in 82.4% of the cases, compared to 74.5% in the CHOICE treatment. Obviously, these comparisons do not account for the increased complexity in Part 3 as well as the repeated nature of decisions.

prosocial, compared to a setting with random matching.

These results have a number of applications, if one is willing to extrapolate to situations outside the lab. First, our setting can be seen as a stylized social media platform, where people follow others based on their information profile. Perhaps disturbingly, the information consumed most frequently on our lab platform consists of irrelevant distractions rather than ethically relevant information. Moreover, even though subjects do not communicate directly about their preferences, our results show similar dynamics of assortative matching as have been found on social media, confirming the “filter bubbles” phenomenon (e.g., [Aiello et al., 2012](#)). As long as people can choose their own connections, they tend to behave in a homophilous way by selecting like-minded sources of information. This also echoes studies on homophily in the endogenous selection of peers in the moral domain, both in the presence ([Gross et al., 2018](#)) and the absence ([Charroin et al., 2021](#)) of complementarity between peers.

The second application is in organizational design, relating to the relationships between an executive and an adviser or consultant. The experiment shows that having decision makers be guided by independent external advisers is not a guarantee for more ethical choices, as advisers may suppress inconvenient information either to please the decision maker or to impose their own agenda in a paternalistic perspective. It also shows that institutional details affect the balance of power. We find that decision makers’ preferences are more predictive of their ethical decisions when they can choose their own advisers. Giving the executive the power to choose advisers puts a greater onus on the executive’s character, while strengthening the independence of the bureaucracy does the opposite, but without leading to more ethical decisions.

Our experiment provides a starting point for further investigations of the transmission of inconvenient information. For instance, as we discussed above, our design suggests that senders’ decisions result from a complex mix of motives. Given how often we make the decision to share (or not) content with others both on and offline, this is an important area for future research. Furthermore, the design of our interactions could be extended in several directions. What happens when sender’s information is known to be more/less noisy? Or when senders can actively falsify, rather than just suppress information, engaging in cheap talk or “fake news” generation rather than disclosure? What if decision makers can consult multiple senders for a second opinion? And what if senders have to bear some accountability for the consequences of the decision makers’ choices? Even if these questions are too numerous to answer in a single paper, our framework could prove useful to investigate them in the future.

References

- Aiello, L. M., A. Barrat, R. Schifanella, C. Cattuto, B. Markines, and F. Menczer (2012): “Friendship prediction and homophily in social media,” *ACM Transactions on the Web (TWEB)*, 6, 9.
- Ambuehl, S., B. D. Bernheim, and A. Ockenfels (2019): “Projective paternalism,” *NBER Working paper 26119*.
- Bartling, B., F. Engl, and R. A. Weber (2014): “Does willful ignorance deflect punishment? An experimental study,” *European Economic Review*, 70, 512–524.
- Bartling, B. and U. Fischbacher (2012): “Shifting the blame: On delegation and responsibility,” *Review of Economic Studies*, 79, 67–87.
- Bartling, B., C. A. W., H. Hermes, M. Skiveness, and B. Tungodden (2020): “Free to Fail: Paternalistic Preferences in the U.S.” *Mimeo obtained via personal interaction*.
- Bartling, B., R. A. Weber, and L. Yao (2015): “Do markets erode social responsibility?” *The Quarterly Journal of Economics*, 130, 219–266.
- Benabou, R., A. Falk, and J. Tirole (2018): “Narratives, Imperatives and Moral Reasoning,” *NBER Working Paper 24798*, July.
- Bock, O., I. Baetge, and A. Nicklisch (2014): “Hroot: Hamburg Registration and Organization Online Tool,” *European Economic Review*, 71, 117–120.
- Charness, G. and M. Sutter (2012): “Groups make better self-interested decisions,” *Journal of Economic Perspectives*, 26, 157–176.
- Charpentier, C. J., E. S. Bromberg-Martin, and T. Sharot (2018): “Valuation of knowledge and ignorance in mesolimbic reward circuitry,” *Proceedings of the National Academy of Sciences of the United States of America*, 115, E7255–E7264.
- Charroin, L., B. Fortin, and M. C. Villeval (2021): “Homophily, Peer Effects and Dishonesty,” *GATE Working Paper*, 2021-07.
- Chen, Z. C. and T. Gesche (2016): “Persistent bias in advice-giving,” *Working Paper*.
- Coffman, L. C. and A. Gotthard Real (2019): “Moral Perceptions of Advised Actions,” *Management Science*, 65, 3904–3927.
- Conrads, J. and B. Irlenbusch (2013): “Strategic ignorance in ultimatum bargaining,” *Journal of Economic Behavior & Organization*, 92, 104–115.
- Cummins, J. G. and I. Nyman (2013): “Yes Men in Tournaments,” *Journal of Institutional and Theoretical Economics*, 169, 621–659.
- Dana, J., D. M. Cain, and R. M. Dawes (2006): “What you dont know wont hurt me: Costly (but quiet) exit in dictator games,” *Organizational Behavior and Human Decision Processes*, 100, 193–201.
- Dana, J., R. A. Weber, and J. X. Kuang (2007): “Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness,” *Economic Theory*, 33, 67–80.

- Di Tella, R., R. Perez-Truglia, A. Babino, and M. Sigman (2015): “Conveniently upset: Avoiding altruism by distorting beliefs about others’ altruism,” *American Economic Review*, 105, 3416–42.
- Ehrich, K. R. and J. R. Irwin (2005): “Willful ignorance in the request for product attribute information,” *Journal of Marketing Research*, XLii, 266–277.
- Espinosa, R. and J. Stoop (2019): “Do People Really Want to Be Informed? Ex-ante Evaluations of Information-Campaign Effectiveness,” *Mimeo*.
- Exley, C. L. (2015): “Excusing selfishness in charitable giving: The role of risk,” *Review of Economic Studies*, 83, 587–628.
- Falk, A. and N. Szech (2013): “Morals and markets,” *Science (New York, N.Y.)*, 340, 707–11.
- Feiler, L. (2014): “Testing Models of Information Avoidance with Binary Choice Dictator Games,” *Journal of Economic Psychology*, 45, 253–267.
- Foerster, M. and J. J. Van der Weele (2018a): “Denial and Alarmism in Collective Action Problems,” *Tinbergen Institute Discussion paper*, 019.
- (2018b): “Persuasion , justification and the communication of social impact,” *Tinbergen Institute Discussion paper*, 067.
- Freddi, E. (2019): “Do People Avoid Morally Relevant Information? Evidence from the Refugee Crisis,” *Review of Economics and Statistics*, In press.
- Garcia, T., S. Massoni, and M. C. Villeval (2020): “Ambiguity and excuse-driven behavior in charitable giving,” *European Economic Review*, 124.
- Gneezy, U., S. Saccardo, M. Serra-Garcia, and R. van Vel dhuizen (2020): “Bribing the Self,” *Games and Economic Behavior*, 120, 311–324.
- Gross, J., M. Leib, T. Offerman, and S. Shalvi (2018): “Ethical Free Riding: When Honest People Find Dishonest Partners,” *Psychological Science*, 29, 1956–1968.
- Grossman, Z. (2014): “Strategic Ignorance and the Robustness of Social Preferences,” *Management Science*, 60, 2659–2665.
- Grossman, Z. and J. J. van der Weele (2017): “Self-Image and Willful Ignorance in Social Decisions,” *Journal of the European Economic Association*, 15, 173–217.
- Haisley, E. C. and R. A. Weber (2010): “Self-serving interpretations of ambiguity in other-regarding behavior,” *Games and Economic Behavior*, 68, 614–625.
- Ismayilov, H. and J. Potters (2013): “Disclosing advisor’s interests neither hurts nor helps,” *Journal of Economic Behavior and Organization*, 93, 314–320.
- Kajackaite, A. (2015): “If I close my eyes, nobody will get hurt: The effect of ignorance on performance in a real-effort experiment,” *Journal of Economic Behavior and Organization*, 116, 518–524.

- Kocher, M. G., S. Schudy, and L. Spantig (2018): “I Lie? We Lie! Why? Experimental Evidence on a Dishonesty Shift in Groups,” *Management Science*, 64, 3971–4470.
- Larson, T. and M. C. Capra (2009): “Exploiting moral wiggle room: Illusory preference for fairness? A comment,” *Judgement and Decision Making*, 4, 467–474.
- Lazear, E. E. P., U. Malmendier, and R. A. Weber (2012): “Sorting in experiments with application to social preferences,” *American Economic Journal: Applied Economics*, 4, 136–163.
- Lind, J. T., K. Nyborg, and A. Pauls (2019): “Save the planet or close your eyes? Testing strategic ignorance in a charity context,” *Ecological Economics*, 161, 9–19.
- Momsen, K. and M. Ohndorf (2019): “Information avoidance, selective exposure, and fake (?) news: A green market experiment,” Tech. rep., Working Papers in Economics and Statistics.
- (2020): “When do people exploit moral wiggle room? An experimental analysis in a market setup,” *Ecological Economics*, In Press.
- Nyborg, K. (2011): “I dont want to hear about it: Rational ignorance among duty-oriented consumers,” *Journal of Economic Behavior & Organization*, 79, 263–274.
- Prendergast, C. (1993): “A theory of ”yes men”,” *American Economic Review*, 83, 757–770.
- Robin, S., A. Rusinowska, and M. C. Villeval (2014): “Ingratiation: Experimental evidence,” *European Economic Review*, 66, 16–38.
- Schram, A. and G. Charness (2015): “Inducing social norms in laboratory allocation choices,” *Management Science*, 61, 1531–1546.
- Serra-Garcia, M. and N. Szech (2019): “The (In)Elasticity of Moral Ignorance,” *Mimeo Karlsruhe Institute for Technology*.
- Sharot, T. and C. R. Sunstein (2020): “How people decide what they want to know,” *Nature Human Behavior*, 4, 14–19.
- van der Weele, J. J. (2013): “Inconvenient Truths : Determinants of Strategic Ignorance in Moral Dilemmas,” *SSRN Electronic Journal*.
- Weisel, O. and S. Shalvi (2015): “The collaborative roots of corruption,” *Proceedings of the National Academy of Sciences*, 112, 10651–10656.

A Appendix: Instructions

These instructions were translated from French.

A.1 Instructions for the RANDOM treatment

Welcome to this experiment. Please switch off your mobile phone and refrain from communication with the other participants throughout the experiment, or we must exclude you from the experiment and from all payments. Please read the instructions carefully. Whenever you have a question, please raise your hand or press the red button on the side of your desk and we will come to your desk and answer to your question in private.

You will receive €5 for showing up on time. You can earn additional earnings based on your decisions and the decisions of other participants. The experiment consists of three parts that can include several periods. At the end of the session, the computer program will randomly select one of these periods, each with equal probability, and we will pay you according to your payoff in the selected period. Thus, you should think carefully when making each decision, as it could be the one that will be paid. At the end of the session, your total earnings will be paid to you in cash in a separate room.

For each participant, the experimenters have prepared a donation of €15 to a charity, GiveDirectly. GiveDirectly transfers money to very poor families in developing countries. This charity is rated as one of the 7 top charities in terms of cost-effectiveness by the charity evaluation site GiveWell, above many traditional charities in the world. Here is an excerpt from the website “GiveDirectly.org” presenting its objectives (we have also distributed a document on the operating mode of GiveDirectly and information from Wikipedia):

“We use mobile payments technology to send your donations to extremely poor families in the developing world in the most capital efficient way currently possible. \$0.91 of your dollar ends up in the hands of the poor. Our model is setting the benchmark for philanthropic efficiency around the world. We strive to promote a new approach to philanthropy that uses constant experimentation and analytical rigor to understand the most impactful ways to achieve positive outcomes.”

During the session, we will show you pictures and testimonials of people who have passed the screening of GiveDirectly, and are potential recipients of the donations in this session. Their pictures and testimonials, translated into French, are taken verbatim from the website “GiveDirectly.org” and they may thus include typos.

The experimenters commit on honor to transfer the donations to GiveDirectly after the experiment. Note that the deontological rules of GATE-Lab do not allow deception of participants by the experimenters. So, all promised donations for the selected period at the end of the session will actually be sent to GiveDirectly. If you want more information about the transfer, please contact an experimenter after the session.

However, as we explain below, your choices may lead to a cancellation of the donation prepared by the experimenters, in which case GiveDirectly will not receive a donation for your participation.

The instructions for the first part follow below. The instructions for the next parts will be distributed after all participants have completed each part.

Part 1

In this part you will make two decisions. In each of these decisions, you are asked to choose between “Option 1” and “Option 2”. Both options affect your own payoffs and the donation to GiveDirectly.

For **Decision one** you will see on your screen before your choice the picture and testimonial of a potential recipient of the donation, who has passed the screening by GiveDirectly, as illustrated in the screenshot below. Choosing “Option 1” will result in €9 for yourself and will *not cancel* the donation of €15 by the experimenters to GiveDirectly. Choosing “Option 2” will result in €15 for yourself, and will *cancel* the donation to GiveDirectly.

In **Decision two** as in decision one, you can choose between “Option 1”, which will result in €9 for yourself and “Option 2”, which will result in €15 for yourself. The difference with the first decision is that the program determines randomly which one of the two options will result in a cancellation of the donation of €15 to Give Directly. With 50 chances out of 100, choosing “Option 2” cancels the donation while “Option 1” confirms the donation, just like in decision 1. With 50 chances out of 100 the situation is reversed, so choosing “Option

1” cancels the donation and “Option 2” confirms the donation. You are not informed which situation is chosen by the program, and the consequences for GiveDirectly are replaced by “???”.

Before choosing between “Option 1” and “Option 2”, you have to choose between two types of information.

You can choose “Recipient”. This means that before choosing between “Option 1” and “Option 2”, you will learn which situation was chosen by the computer, and the “???” will be replaced with information about the consequences for GiveDirectly. Furthermore, like in decision 1, your screen will display the picture and testimonial of a potential recipient before your choice of option.

Or you can choose “Cute animal”. This means that your screen will display the picture of a cute animal, as illustrated in the screenshot below. You will not learn which situation was selected by the computer, neither before nor after your choice of option.

After making this choice, we will inform you about the number of participants in this session and ask you to guess the number of participants who have chosen “Recipient” and the number of participants who have chosen “Cute animal”. Regardless of whether this period is selected or not for payment at the end of the session, you will earn 1 euro if your guess is correct, and 0 euro otherwise. Therefore, you should try to guess as accurately as possible. You will be informed on whether your guess is correct at the end of the session.

After you have chosen between “Recipient” and “Cute animal” and reported your guess, you will have to choose between “Option 1” and “Option 2”. Your earlier choice between “Recipient” or “Cute animal” determines the information you see on your screen before making your choice.

Summary of the Decisions

1. In Decision 1, you choose between Option 1 and Option 2.
2. In Decision 2, the program randomly selects which one of the two options cancels the donation.
3. You choose between the sets of information “Recipient” or “Cute animal”.
4. You report your guess about the numbers of other participants in the session who chose “Recipient” or “Cute animal”.
5. Your screen displays the information you chose in step 3, and you choose between Option 1 and Option 2.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

Part 2

In this part, you are randomly matched with 6 other participants to form a group of 7. There are two roles: Receivers and Senders. Receivers and Senders refer not to donations but to pictures and information, as explained below. All the participants in the group will first make decisions as Senders. Then, all of them will make a decision as Receivers. We first describe each role before explaining decision-making.

Choice of the Receiver

The Receiver has to choose between “Option 1” and “Option 2”. The consequences from this decision are the same as in the second decision of part 1:

“Option 1” results in €9 for the Receiver and “Option 2” results in €15 for the Receiver.

The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50 chances out of 100 to be picked.

The program randomly determines the consequences of each option independently for each Receiver. Thus, these consequences can differ across Receivers. Before making a choice between Option 1 and Option 2, the Receiver is not informed of the consequences of each option for GiveDirectly. However, s/he can obtain information from the Sender, as we now describe.

Choice of the Sender

With 80% chance, the Sender learns which one of the two options cancels the donation. With 20% chance the Sender does not learn the consequences of each option.

If the Sender does not learn the consequences of each option for GiveDirectly, the program displays automatically the picture of a cute animal on the Receiver’s screen before s/he makes his/her choice. The Receiver is not informed on the consequences of this option for GiveDirectly.

If the Sender learns the consequences of each option for GiveDirectly, s/he has to choose between two types of information for the Receiver. If s/he chooses “Recipient”, the Receiver will learn which one of the two options cancels the donation before choosing an option, and s/he will see the picture and the testimonial of a potential recipient of the donation. If the Sender chooses “Cute animal”, the Receiver will see the picture of a cute animal, but not the consequences for GiveDirectly, neither before nor after the choice of option.

Decision Making in Part 2

In this part, all the participants first make two decisions in the role of a Sender in the case they are informed about the consequences of each option for GiveDirectly. Precisely, as a Sender, you have to choose between two types of information for the Receiver, either “Recipient” or ‘Cute animal”:

in the case you learn that the donation to GiveDirectly is cancelled after “Option 1” , but not after “Option 2”;

and in the case you learn that the donation to GiveDirectly is cancelled after “Option 2”, but not after “Option 1”.

Then, all the participants will make a decision as Receivers. As a Receiver, you will have to choose between “Option 1” and “Option 2”. Before you make your choice, the computer will randomly determine which option cancels the donation. It will also randomly pair you with a Sender in your group. The choice of the Sender between “Recipient” or ‘Cute animal” determines the information you have about the consequences of each option.

Summary of the Decisions

1. You first decide as a Sender which picture and information to share if you are informed of the consequences of each option for Give Directly.
2. You are next a Receiver. You are randomly matched with a Sender.
3. You obtain the information chosen by the sender, “Recipient” or ‘Cute animal”.
4. You choose between Option 1 and Option 2.
5. You are paid based on your choice as a Receiver in case this part is selected for payment.

Please read again these instructions. If you have any questions, please raise your hand or press the red button.

Part 3

In this part, you are still matched with the same 6 other participants as in part 2. But now, participants are randomly assigned to one of the roles and will be identified with an ID. There are four Receivers and their IDs are R1, R2, R3, and R4. There are three Senders and their IDs are symbols (spade, diamond, club). We will communicate your role and your ID on your screen at the beginning of this part. This part has 25 identical periods and you will keep the same role and the same ID throughout this part. We now describe each of these periods.

Choice of the Receiver

In each period, the Receiver chooses between “Option 1” and “Option 2”. The consequences from this decision are the same as before:

“Option 1” results in €9 for the Receiver and “Option 2” results in €15 for the Receiver.

The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50% chance to be picked.

In each period, the program randomly determines the consequences of each option for GiveDirectly, independently for each Receiver. Thus, these consequences can differ across periods and across Receivers.

The Receiver is not informed about the consequences of each option for GiveDirectly. Before choosing between Option 1 and Option 2, s/he can receive information from the Sender.

Choice of the Sender

With 80% chance, the Sender learns the consequences of each option for GiveDirectly chosen by the program for each Receiver for the current period. If the Sender learns the consequences, s/he has to decide which set of information to share with the Receiver. As before, if s/he chooses “Recipient”, i) the Receiver is informed which option cancels the donation before choosing an option, and ii) the picture of a potential recipient with his/her testimonial is displayed. If the Sender chooses “cute animal” the picture of a cute animal is displayed on the Receiver’s screen and the Receiver does not learn which option cancels the donation, neither before nor after the

choice of option.

With 20% chance the Sender does not learn the consequences of each option and the picture of a cute animal is displayed automatically on the Receiver's screen. The Receiver does not know whether the Sender has been informed or not.

The decision of the Sender is illustrated in the screenshot below. The first column of the table shows the ID of the Receivers (i.e. R1, R2, R3 or R4). The second column indicates for each Receiver, which option cancels the donation. The "???" sign indicates that the Sender did not receive information for participant R3; in this case, the sign of a cute animal is automatically pre-selected with no action from the Sender. In the next column, when informed, the Sender has to choose between "Recipient" and "Cute animal". In this example, the Sender makes three decisions, as s/he has information about the consequences of each option for three out of four Receivers. Once the Sender has made his/her decisions, the last column of the table indicates which information will be displayed on the screen of the Receiver.

As we explain now, the Sender can earn €10 for each Receiver to whom s/he has been randomly matched by the program in that period.

The Receiver is Matched with a Sender

Before the Receiver chooses an option, s/he is randomly matched with one of the Senders (spade, diamond, or club) for the current period. The screenshot below reflects the screen the Receiver will see in the experiment. The example shows period 5. The first column shows the ID of each Sender. For each of the past periods, the screen shows which set of information each Sender shared with the Receiver (indicated by the symbol of an animal or the symbol GD for a recipient of GiveDirectly). The symbol of an animal reflects *either* that the Sender had no information, *or* that the Sender received the information and decided to share this picture; the Receiver cannot distinguish between these possibilities. The past choices of the Senders in the group are only visible to the Receivers, not to the other Senders. In this example, symbols have been chosen randomly.

Before being randomly matched by the program to a Sender, the Receiver only knows the past choices of the Senders but not those for the current period. The Receiver's information ("Recipient" or "Cute animal") is determined by the randomly matched sender's choice for the current period. The Senders are not informed of the choices of the Receivers.

Summary of the Decisions

1. For each Receiver, the program randomly selects which option cancels the donation.
2. With 80% chance, the Sender learns the consequences of each option. If s/he learns the consequences, the Sender has to choose a set of information to share with the Receiver (“Recipient” or “Cute animal”).
3. The Receiver sees information about the senders’ choices in earlier periods. S/he is randomly matched by the program to one of the Senders. The Receiver’s information is determined by the choice of that sender (“Recipient” or “Cute animal”) in the current period.
4. The Receiver decides between Option 1 and Option 2.
5. The same procedure applies for the 25 periods.

Summary of Earnings

The program randomly selects one of the 28 periods for payment (the two decisions in part 1, the decision as a Receiver in part 2 and the 25 periods in part 3). If a decision in part 1 or in part 2 is selected for payment, you will receive either €9 or €15, depending on whether you chose option Option 1 or Option 2 in that period. If a decision in part 3 is selected, payoffs for each player are as follows:

The Receiver earns either €9 or €15, depending on the chosen option in that period.

The Sender earns €10 for each Receiver to whom s/he has been randomly matched by the program in that period. Thus, the Sender minimally earns €0 if s/he has not been matched to any Receiver, and maximally earns €40 if s/he has been matched to the four Receivers.

Finally, **GiveDirectly** will receive a donation of €15 for any choice in the selected period that does not cancel the donation.

End of the Session

At the end of part 3 a questionnaire will be displayed on your screen and then you will receive a feedback on your earnings in the session. On invitation of an experimenter, you will move into the payment room with your pre-filled receipt of payment and your computer tag.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

A.2 Instructions for the CHOICE treatment

Welcome to this experiment. Please switch off your mobile phone and refrain from communication with the other participants throughout the experiment, or we must exclude you from the experiment and from all payments. Please read the instructions carefully. Whenever you have a question, please raise your hand or press the red button on the side of your desk and we will come to your desk and answer to your question in private.

You will receive €5 for showing up on time. You can earn additional earnings based on your decisions and the decisions of other participants. The experiment consists of three parts that can include several periods. At the end of the session, the computer program will randomly select one of these periods, each with equal probability, and we will pay you according to your payoff in the selected period. Thus, you should think carefully when making each decision, as it could be the one that will be paid. At the end of the session, your total earnings will be paid to you in cash in a separate room.

For each participant, the experimenters have prepared a donation of €15 to a charity, GiveDirectly. GiveDirectly transfers money to very poor families in developing countries. This charity is rated as one of the 7 top charities in terms of cost-effectiveness by the charity evaluation site GiveWell, above many traditional charities in the world. Here is an excerpt from the website “GiveDirectly.org” presenting its objectives (we have also distributed a document on the operating mode of GiveDirectly and information from Wikipedia):

We use mobile payments technology to send your donations to extremely poor families in the developing world in the most capital efficient way currently possible. \$0.91 of your dollar ends up in the hands of the poor. Our model is setting the benchmark for philanthropic efficiency around the world. We strive to promote a new approach to philanthropy that uses constant experimentation and analytical rigor to understand the most impactful ways to achieve positive outcomes.

During the session, we will show you pictures and testimonials of people who have passed the screening of GiveDirectly, and are potential recipients of the donations in this session. Their pictures and testimonials, translated into French, are taken verbatim from the website “GiveDirectly.org” and they may thus include typos.

The experimenters commit on honor to transfer the donations to GiveDirectly after the experiment. Note that the deontological rules of GATE-Lab do not allow deception of participants by the experimenters. So, all promised donations for the selected period at the end of the session will actually be sent to GiveDirectly. If you want more information about the transfer, please contact an experimenter after the session.

However, as we explain below, your choices may lead to a cancellation of the donation prepared by the experimenters, in which case GiveDirectly will not receive a donation for your participation.

The instructions for the first part follow below. The instructions for the next parts will be distributed after all participants have completed each part.

Part 1

In this part you will make two decisions. In each of these decisions, you are asked to choose between “Option 1” and “Option 2”. Both options affect your own payoffs and the donation to GiveDirectly.

For **Decision one** you will see on your screen before your choice the picture and testimonial of a potential recipient of the donation, who has passed the screening by GiveDirectly, as illustrated in the screenshot below. Choosing “Option 1” will result in €9 for yourself and will *not cancel* the donation of €15 by the experimenters to GiveDirectly. Choosing “Option 2” will result in €15 for yourself, and will *cancel* the donation to GiveDirectly.

In **Decision two** as in decision one, you can choose between “Option 1”, which will result in €9 for yourself and “Option 2”, which will result in €15 for yourself. The difference with the first decision is that the program determines randomly which one of the two options will result in a cancellation of the donation of €15 to GiveDirectly. With 50 chances out of 100, choosing “Option 2” cancels the donation while “Option 1” confirms the donation, just like in decision 1. With 50 chances out of 100 the situation is reversed, so choosing “Option

1” cancels the donation and “Option 2” confirms the donation. You are not informed which situation is chosen by the program, and the consequences for GiveDirectly are replaced by “???”.

Before choosing between “Option 1” and “Option 2”, you have to choose between two types of information.

You can choose “Recipient”. This means that before choosing between “Option 1” and “Option 2”, you will learn which situation was chosen by the computer, and the “???” will be replaced with information about the consequences for GiveDirectly. Furthermore, like in decision 1, your screen will display the picture and testimonial of a potential recipient before your choice of option.

Or you can choose “Cute animal”. This means that your screen will display the picture of a cute animal, as illustrated in the screenshot below. You will not learn which situation was selected by the computer, neither before nor after your choice of option.

After making this choice, we will inform you about the number of participants in this session and ask you to guess the number of participants who have chosen “Recipient” and the number of participants who have chosen “Cute animal”. Regardless of whether this period is selected or not for payment at the end of the session, you will earn 1 euro if your guess is correct, and 0 euro otherwise. Therefore, you should try to guess as accurately as possible. You will be informed on whether your guess is correct at the end of the session.

After you have chosen between “Recipient” and “Cute animal” and reported your guess, you will have to choose between “Option 1” and “Option 2”. Your earlier choice between “Recipient” or “Cute animal” determines the information you see on your screen before making your choice.

Summary of the Decisions

1. In Decision 1, you choose between Option 1 and Option 2.
2. In Decision 2, the program randomly selects which one of the two options cancels the donation.
3. You choose between the sets of information “Recipient” or “Cute animal”.
4. You report your guess about the numbers of other participants in the session who chose “Recipient” or “Cute animal”.
5. Your screen displays the information you chose in step 3, and you choose between Option 1 and Option 2.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

Part 2

In this part, you are randomly matched with 6 other participants to form a group of 7. There are two roles: Receivers and Senders. Receivers and Senders refer not to donations but to pictures and information, as explained below. All the participants in the group will first make decisions as Senders. Then, all of them will make a decision as Receivers. We first describe each role before explaining decision-making.

Choice of the Receiver

The Receiver has to choose between “Option 1” and “Option 2”. The consequences from this decision are the same as in the second decision of part 1:

“Option 1” results in €9 for the Receiver and “Option 2” results in €15 for the Receiver.

The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50 chances out of 100 to be picked.

The program randomly determines the consequences of each option independently for each Receiver. Thus, these consequences can differ across Receivers. Before making a choice between Option 1 and Option 2, the Receiver is not informed of the consequences of each option for GiveDirectly. However, s/he can obtain information from the Sender, as we now describe.

Choice of the Sender

With 80% chance, the Sender learns which one of the two options cancels the donation. With 20% chance the Sender does not learn the consequences of each option.

If the Sender does not learn the consequences of each option for GiveDirectly, the program displays automatically the picture of a cute animal on the Receiver’s screen before s/he makes his/her choice. The Receiver is not informed on the consequences of this option for GiveDirectly.

If the Sender learns the consequences of each option for GiveDirectly, s/he has to choose between two types of information for the Receiver. If s/he chooses “Recipient”, the Receiver will learn which one of the two options cancels the donation before choosing an option, and s/he will see the picture and the testimonial of a potential recipient of the donation. If the Sender chooses “Cute animal”, the Receiver will see the picture of a cute animal, but not the consequences for GiveDirectly, neither before nor after the choice of option.

Decision Making in Part 2

In this part, all the participants first make two decisions in the role of a Sender in the case they are informed about the consequences of each option for GiveDirectly. Precisely, as a Sender, you have to choose between two types of information for the Receiver, either "Recipient" or "Cute animal":

in the case you learn that the donation to GiveDirectly is cancelled after "Option 1", but not after "Option 2";

and in the case you learn that the donation to GiveDirectly is cancelled after "Option 2", but not after "Option 1".

Then, all the participants will make a decision as Receivers. As a Receiver, you will have to choose between "Option 1" and "Option 2". Before you make your choice, the computer will randomly determine which option cancels the donation. It will also randomly pair you with a Sender in your group. The choice of the Sender between "Recipient" or "Cute animal" determines the information you have about the consequences of each option.

Summary of the Decisions

1. You first decide as a Sender which picture and information to share if you are informed of the consequences of each option for Give Directly.
2. You are next a Receiver. You are randomly matched with a Sender.
3. You obtain the information chosen by this sender, "Recipient" or "Cute animal".
4. You choose between Option 1 and Option 2.
5. You are paid based on your choice as a Receiver in case this part is selected for payment.

Please read again these instructions. If you have any questions, please raise your hand or press the red button.

Part 3

In this part, you are still matched with the same 6 other participants as in part 2. But now, participants are randomly assigned to one of the roles and will be identified with an ID. There are four Receivers and their IDs are R1, R2, R3, and R4. There are three Senders and their IDs are symbols (spade, diamond, club). We will communicate your role and your ID on your screen at the beginning of this part. This part has 25 identical periods and you will keep the same role and the same ID throughout this part. We now describe each of these periods.

Choice of the Receiver

In each period, the Receiver chooses between "Option 1" and "Option 2". The consequences from this decision are the same as before:

"Option 1" results in €9 for the Receiver and "Option 2" results in €15 for the Receiver.

The program picks randomly which one of the two options cancels the donation to GiveDirectly. Each option has 50% chance to be picked.

In each period, the program randomly determines the consequences of each option for GiveDirectly, independently for each Receiver. Thus, these consequences can differ across periods and across Receivers.

The Receiver is not informed about the consequences of each option for GiveDirectly. Before choosing between Option 1 and Option 2, s/he can receive information from the Sender.

Choice of the Sender

With 80% chance, the Sender learns the consequences of each option for GiveDirectly chosen by the program for each Receiver for the current period. If the Sender learns the consequences, s/he has to decide which set of information to share with the Receiver. As before, if s/he chooses "Recipient", i) the Receiver is informed which option cancels the donation before choosing an option, and ii) the picture of a potential recipient with his/her testimonial is displayed. If the Sender chooses "cute animal" the picture of a cute animal is displayed on the Receiver's screen and the Receiver does not learn which option cancels the donation, neither before nor after the

choice of option.

With 20% chance the Sender does not learn the consequences of each option and the picture of a cute animal is displayed automatically on the Receiver's screen. The Receiver does not know whether the Sender has been informed or not.

The decision of the Sender is illustrated in the screenshot below. The first column of the table shows the ID of the Receivers (i.e. R1, R2, R3 or R4). The second column indicates for each Receiver, which option cancels the donation. The "???" sign indicates that the Sender did not receive information for participant R3; in this case, the sign of a cute animal is automatically pre-selected with no action from the Sender. In the next column, when informed, the Sender has to choose between "Recipient" and "Cute animal". In this example, the Sender makes three decisions, as s/he has information about the consequences of each option for three out of four Receivers. Once the Sender has made his/her decisions, the last column of the table indicates which information will be displayed on the screen of the Receiver.

As we explain now, the Sender can earn 10 for each Receiver that selects him/her in that period.

The Receiver Selects a Sender

Before the Receiver chooses an option, s/he has to select one of the Senders (spade, diamond, or club) for the current period. The selection decision is illustrated in the screenshot below, which reflects the screen the Receiver will see in the experiment. The example shows the decision in period 5. The first column shows the ID of each Sender. For each of the past periods, the screen shows which set of information each Sender shared with the Receiver (indicated by the symbol of an animal or the symbol GD for a recipient of GiveDirectly). The symbol of an animal reflects either that the Sender had no information, or that the Sender received the information and decided to share this picture; the Receiver cannot distinguish between these possibilities. The past choices of the Senders in the group are only visible to the Receivers, not to the other Senders. In this example, symbols have been chosen randomly.

When choosing a Sender, the Receiver only knows the past choices of the Senders but not those for the current period. The Receiver chooses a Sender by clicking on a box in the left column. The Receiver's information ("Recipient" or "Cute Animal") is determined by the selected sender's choice for the current period. The Senders are not informed of the choices of the Receivers.

Summary of the Decisions

1. For each Receiver, the program randomly selects which option cancels the donation.

2. With 80% chance, the Sender learns the consequences of each option. If s/he learns the consequences, the Sender has to choose a set of information to share with the Receiver ("Recipient" or "Cute animal").
3. The Receiver sees information about sender's choices in earlier periods. S/he then chooses one of the Senders. The Receiver's information is determined by the choice of that sender ("Recipient" or "Cute animal") in the current period.
4. The Receiver decides between Option 1 and Option 2.
5. The same procedure applies for the 25 periods.

Summary of Earnings

The program randomly selects one of the 28 periods for payment (the two decisions in part 1, the decision as a Receiver in part 2 and the 25 periods in part 3). If a decision in part 1 or in part 2 is selected for payment, you will receive either e9 or e15, depending on whether you chose option Option 1 or Option 2 in that period. If a decision in part 3 is selected, payoffs for each player are as follows:

The Receiver earns either e9 or e15, depending on the chosen option in that period.

The Sender earns e10 for each Receiver that selected him/her in that round. Thus, the Sender minimally earns e0 if s/he has not been chosen by any Receiver, and maximally earns e40 if s/he has been chosen by the four Receivers.

Finally, GiveDirectly will receive a donation of e15 for any choice in the selected period that does not cancel the donation.

End of the Session

At the end of part 3 a questionnaire will be displayed on your screen and then you will receive a feedback on your earnings in the session. On invitation of an experimenter, you will move into the payment room with your pre-filled receipt of payment and your computer tag.

Please read again these instructions. If you have any questions, please raise your hand or press the red button. A comprehension questionnaire will be displayed on your screen.

B Appendix: Additional Tables and Figures

Summary of the sessions

Table B.1: Summary of Sessions

Session	Treatment	N Participants	Females	Mean Age	Mean Payo
1	CHOICE	21	66.67%	29.09	18.76
2	CHOICE	21	61.09%	23.38	18.52
3	CHOICE	21	52.38%	23.67	18.81
4	CHOICE	21	52.38%	23.43	19.33
5	CHOICE	21	50.00%	21.45	18.33
6	CHOICE	21	57.14%	20.29	18.76
7	CHOICE	21	42.86%	21.47	19.05
8	CHOICE	14	40.00%	20.87	18.21
9	RANDOM	21	42.86%	21.48	18.86
10	RANDOM	21	80.95%	20.76	17.67
11	RANDOM	21	57.14%	20.62	18.48
12	RANDOM	21	57.14%	23.52	15.48
13	RANDOM	21	61.90%	23.43	18.24
14	RANDOM	21	61.90%	21.68	18.90
15	RANDOM	14	57.14%	20.79	19.57
16	RANDOM	21	38.10%	23.10	19.14
Total	-	322	55.28%	22.50	18.49

Notes: The table reports the number of participants, the percentage of males, the mean age of the participants, and the mean participant's payo in Euros, per session. The smaller number of participants in two sessions (one per treatment) is due to no show-up. The high mean age in session 1 is due to the presence of two participants aged 60 and 63.

C Appendix: Additional Figures on Supply Choices

In this Appendix, we show more details of the distribution of supply, using our metric for the suppression of bad news, which we call the s statistic, defined as

$$s_i := \frac{\text{Number of bad states suppressed by sender } i}{\text{Number of bad states observed by sender } i}$$

We also provide additional figures displaying the roles of senders' beliefs and senders' suppression of good news.

Supply: Overview of Suppression of Bad News.

Figure C.1 shows the cumulative distribution of s -statistics over senders in each treatment, which reveals several results. First, the distributions do not differ much by treatment. Indeed, a Kolmogorov-Smirnov test cannot reject equality of the distributions ($p = 0.248$). Second, almost one third of senders in each treatment transmits all bad news and the large majority suppresses at least some news. Third, about 25% of senders suppress more than half of the bad news they receive ($s > 0.5$). Appendix C gives individual examples to illustrate various patterns of information suppression by senders with different s -statistics.

Figure C.1: Suppression of bad news by senders.

Note: The figure displays the cumulative distribution of individual s -statistics by treatment.

Supply: Correlation between beliefs and suppression of bad news.

The left panel of Figure C.2 shows the suppression rates by senders who believe that less than 50% of other participants in the session prefer animal pictures over recipients' info ("belief low") and senders who believe the fraction is 50% or more ("belief high"). In both treatments, the (minority of) senders who believe the majority prefers distractions are more likely to suppress information. However, contrary to a strategic motive of attracting clients, this difference is larger in the RANDOM treatment than in the CHOICE treatment.

These results are consistent with the idea that senders are trying to do decision makers a "favor", rather than strategically adjust their behavior to increase their chance of being selected by decision makers. However, beliefs may also reflect a "false consensus effect" and be a proxy for senders' own preferences for information. If so, our findings could indicate "paternalism": a wish to impose the sender's preferred information or decision on the decision maker.

To investigate this possibility, we consider the correlation between sender beliefs and sender preferences. The right panel of Figure C.2 shows the beliefs split by the senders' own preferences in the dilemma, based on the individual decisions in Part 1. In line with the "projection" hypothesis, senders who prefer to avoid receiver information ("Avoiders") are more likely to think that others prefer to do so. The impact of taking a selfish decision in the dilemma ("Selfish") is less clear.

Figure C.2: Sender beliefs and suppression of bad news

Notes: The left panel shows the impact of sender's beliefs about the number of subjects (in the session) that want to remain uninformed on the rate of suppression of bad news. Belief high (low) is for senders that believe the majority wants to remain ignorant (wants to know). The right panel shows the average belief about the number of subjects that want to remain ignorant conditional on the sender's preferences as revealed in Part 1 of the experiment. Vertical bars show the standard errors. In the left panel, standard errors are based on a linear probability model that clusters errors at group level and it is analogous to the ones reported in Table 1. Compared to the table, the model only includes the treatment dummy, a dummy that is equal to 1 when the sender believes that the majority wants to remain uninformed, and their interaction as explanatory variables.

Supply: Suppression of Good News.

Although the suppression of good news does not affect the decision makers' choice, senders may suppress information symmetrically to avoid that decision makers infer that "no news means bad news". To investigate this, Figure C.3 shows each sender's suppression pattern for both good and bad news. We observe that most of the observations are close to the diagonal, indicating symmetric suppression. Fisher tests reject the null hypothesis that the senders have the same fraction of suppression for bad and good news for only 35 out of 138 senders. Of these, 21 systematically suppress bad news more than good news and 14 suppress good news more than bad news. Thus, only a minority systematically suppress news asymmetrically, which is consistent with the idea that senders think about the (Bayesian) inferences of the decision maker.

At first glance, it is somewhat puzzling that 14 senders suppress good news more than bad news. One rationale for suppressing mostly good news is to signal information suppression in order to lure decision makers who would prefer to remain ignorant. Indeed, such a strategy makes the picture of the beneficiary (and potential victim) more salient when news are bad. For example, one subject reported in the final questionnaire: "When option 1 cancelled the donation, regardless of whether I give information or not, the Receiver would choose to earn 15 Euros (option2), which is good since it gives money to the charity. Even if it has no impact on their decision, I prefer send them a picture of a cute animal. When option 2 cancelled the donation, I sent the picture of GD hoping that this would force the receivers to think further, so that the charity would receive a donation (option 1), even if they would earn only 9 Euros." By providing them with bad news, an altruistic sender may exert pressure on such decision makers to act altruistically, thus increasing revenue for the charity. This idea is consistent with the data: of the 14 decision makers who mostly suppress good news, 12 are Altruistic and 2 are Selfish. Again, this suggests that senders employ rather sophisticated strategies.

Figure C.3: Individual suppression rates of good and bad news by treatment

Notes: The figure displays suppression of good news (x-axis) and bad news (y-axis). Each dot is a sender. Multiple senders with the same suppression patterns are indicated by bigger dots.

Supply: Illustrations of the supply of information by sender with various s-statistics

The following three figures correspond to three senders with different s-statistics. In the figures, each line corresponds to one of the four decision makers in the sender's group. The horizontal axis indicates the 25 periods in part 3 of the experiment. B is for bad news (option 2 cancels the donation); G for good news (option 1 cancels the donation); and an hyphen indicates that no news has been sent (either because the sender was not informed or because he or she decided not to send news). Colored letters indicate that the corresponding news has been sent and dark letters that the news has not been sent.

Figure C.4: Transmission of information to decision makers by a sender with s -statistic = 0

Notes: Figure C.4 illustrates the case of a sender that transmits all information (s -statistic = 0).

Figure C.5: Transmission of information to decision makers by a sender with s -statistic = 0.66

Notes: Figure C.5 illustrates the case of a sender that suppresses all types of news except for one decision maker (s -statistic = 0.66).

Figure C.6: Transmission of information to decision makers by a sender with s -statistic = 0.89

Notes: Figure C.6 illustrates the case of a sender that suppresses most bad news and does not discriminate among decision makers (s -statistic = 0.89). This sender sends news almost only when they are good.

D Appendix: Additional Analysis of the Demand for Information

In this Appendix, we provide additional figures and tables to understand information demand, and use a number of additional methods to classify decision makers' selection strategies.

Probability to choose the senders | Multinomial models

Table D.1 reports multinomial logit models where the three alternatives are the senders providing Low, Medium, and High ignorance. The individual specific explanatory variables include the dummies capturing the type of the decision maker, i.e., *Sel sh* - Altruistic and *Avoider* - Seeker, obtained from the choices in Part 1. Both Model 1 and Model 2 include random effects at the individual level on the intercepts of the two equations. Compared to Model 1, Model 2 includes fixed effects at the group level using group dummies.

The estimates in Table D.1 show the effect of the decision maker types on the odds ratio of Low versus Medium and of Low versus High ignorance, respectively. Moving from Altruistic to *Sel sh* and from information Seeking to Avoiding significantly lowers the odds to choose the sender that provides the lowest level of ignorance compared to the odds to choose the sender that provides the intermediate level (Medium ignorance equation), and compared to the odds to choose the sender that provides the highest level of ignorance (High ignorance equation). Results are robust to the inclusion of fixed effects at the group level.

Table D.1: Probability to choose the sender that provided the highest, intermediate and lowest level of ignorance.

	Model 1 Est. (S.E.)	Model 2 Est. (S.E.)
Medium ignorance		
(intercept)	-1.010 (0.109)***	-0.000 (0.379)
<i>Sel sh</i> (DM)	0.349 (0.181)	0.476 (0.200)*
<i>Avoider</i> (DM)	1.013 (0.247)***	1.089 (0.282)***
<i>Sel sh</i> (DM) <i>Avoider</i> (DM)	-0.675 (0.365)	-0.844 (0.430)*
<i>M</i>	0.779 (0.105)***	0.498 (0.127)***
High ignorance		
(intercept)	-1.037 (0.116)***	0.355 (0.405)
<i>Sel sh</i> (DM)	0.848 (0.188)***	0.741 (8.218)***
<i>Avoider</i> (DM)	1.573 (0.253)***	1.647 (0.301)***
<i>Sel sh</i> (DM) <i>Avoider</i> (DM)	0.001 (0.385)	-0.917 (0.441)*
<i>H</i>	1.560 (0.140)***	1.190 (0.139)***
Group dummies	NO	YES
Log-Likelihood	-1327.3	-1293.4
Number of observations	1380	1380
Number of subjects	92	92
Number of groups	23	23

Notes: These regressions are based on a multinomial logit model where the alternatives are the three senders ordered by the amount of ignorance supplied in the previous 10 periods (the baseline alternative is the sender that supplies the lowest level of ignorance). Individual specific variables are the dummies indicating the preferences of the decision makers. Both models include random effects at subject level on the intercepts. Model 2 includes group dummies. Regressions use data of the last 15 periods. *** 0.001; ** 0.01; * 0.05; 0.1.

To help interpret these findings, Figure D.1 shows the predicted probability to choose each sender for each decision maker. It shows that the effect of heterogeneity is mostly captured by shifting the probability mass from the Low ignorance to the High ignorance sender, while the predicted probability to choose the Medium ignorance sender is about 20-25% and does not change much across types. Both the *Sel sh* and Altruistic Avoiders show

a significantly lower propensity to choose informative senders.

Figure D.1: Predicted probability of the decision makers' choice of sender

Notes: The figures display the predicted probability to choose the sender providing the highest, intermediate, and lowest level of suppression for each of the 92 decision makers. The figure on the left reports predictions based on Model 1 of Table D.1 and the figure on the right reports predictions based on Model 2 of Table D.1. The color of the dots captures the type of the decision maker elicited in Part 1. The cross shows the average frequency of choice.

Demand as measured by news consumption

Here, we compute how often a decision maker actually observes bad news, compared to the bad news that is available from senders in the group. We focus on bad news, since this is the only news that matters from an ethical or efficiency perspective. We define the d -statistic, which is the fraction of bad states seen by the decision maker out of the average number of bad states reported to the decision maker (DM) by the three senders he or she was matched with throughout the 25 periods of Part 3:

$$d_i := \frac{\text{Number of bad states seen by DM } i}{\frac{1}{3} \sum_{j=1}^3 \text{Bad states reported to the DM } i \text{ by sender } j}.$$

Selecting a sender at random will lead to $d = 1$. A decision maker who consistently selects informative senders will have $d > 1$, whereas selecting uninformative senders will yield $d < 1$.

Figure D.2: Information demand

Notes: The figure displays the distribution of the d -statistic across subjects in the CHOICE treatment. A d -statistic equal to 1 corresponds to a random choice of senders; a d -statistic higher (lower) than 1 corresponds to the selection of informative (uninformative) senders. The left panel ranks the individual d -statistics by size (dots and triangles indicate the observed d -statistic) and shows a 90% confidence interval for the subject's d -statistic under the null hypothesis of random sender choice (black bars). When the observed d -statistic lies outside the black bars (which is indicated with a triangle), one can reject the null hypothesis of a random selection of sender at the 10% confidence level. The right panel shows the overall distribution of the observed d -statistic.

Figure D.2 shows the distribution of the d -statistic. The left panel ranks all individual d -statistics by size, whereas the right panel shows the density distribution of the d -statistic. The left panel tells us that 40 out of 92 (43.5%) decision makers have $d < 1$, and can be classified as information avoiders in Part 3, while the rest consists of information seekers. Appendix C illustrates the demand for information by decision makers with different d -statistics. However, the d -statistic is a noisy measure since groups differ in the distribution of news, and hence in the possibility to become more or less informed. In an extreme case where all senders transmit the same amount of news, the d -statistic will necessarily be 1, no matter what the news consumption is. We control for this by conducting simulations based on a decision maker who chooses randomly. This yields a distribution of d -statistics that we use to construct a 90% confidence interval.¹⁶ By comparing the actual d -statistic to this

¹⁶The procedure is as follows: (i) keeping fixed the senders' behavior, we simulate the choice of each decision maker in each period under the assumption that he/she randomly selects one of the 3 senders; (ii) given the simulated choices of the decision maker, we compute the implied d -statistic; (iii) we repeat the procedure 100000 times. The confidence intervals are obtained by taking the 5th and 95th percentiles of the simulated d -statistic over the 100000 simulations. This interval captures the most likely values of the d -statistic under the assumption that the decision maker is neutral to the information received.

confidence interval, we can classify with 90% confidence 13 decision makers as information avoiders (identifiable by triangles on the left hand side of the left panel), and 21 as information seekers (identifiable by triangles on the right hand side of the left panel), out of a total of 92. The percentage of avoiders in the set of clearly classifiable decision makers is 38.2%, not too far from the 43.5% we found above.

Overall, this analysis confirms that a sizable minority of subjects appears to either avoid informative senders or at least not seek them out. Coming back to our type classification of decision makers, the d -statistic correlates with the decision to avoid information in the first individual part of the experiment (Pearson $r = 0.202$, $p = 0.054$).

Illustrations of the demand for information by decision makers with various d -statistics

The following three figures correspond to three decision makers with different d -statistics. In the figures, each line corresponds to one of the three senders in the decision maker's group. The horizontal axis indicates the 25 periods in part 3 of the experiment. B is for bad news (option 2 cancels the donation); G for good news (option 1 cancels the donation); and an hyphen indicates that no news has been received. The colored items indicate which advisor has been selected in each period and which news has been revealed in the period after the sender has been selected.

Figure D.3: Demand for information of a decision maker with d -statistic = 1.23

Note: Figure D.3 illustrates the case of an information seeker who sanctions any transmission of no news (d -statistic = 1.23).

Figure D.4: Demand for information of a decision maker with d -statistic = 0.75

Notes: Figure D.4 illustrates also the case of an information avoider who is less able to establish a stable relationship with a sender (d -statistic = 0.75).

Figure D.5: Demand for information of a decision maker with d -statistic = 0.18

Notes: Figure D.5 illustrates the case of an information avoider who establishes a long term relationship with a sender who never provides news (d -statistic = 0.18).

Decision to switch sender | Linear probability models

Here we present a third way to analyze demand, namely by the decision to switch to another sender conditional on the information received. This decision is informative of decision makers' information seeking or avoiding strategy. It provides an alternative way to look at information demand based on sender history. Figure D.6 shows the fraction of decision makers that change sender in part 3 in the CHOICE treatment after sender reported good, bad or no news. The left panel shows aggregate results, which demonstrates that switching rates are substantial and vary between 47.3% after bad news, 43.1% after good news, and 56.9% after no news. On aggregate, switching is highest after no news, in line with the idea that most people are information seekers.

Figure D.6: Switching senders

Notes: The figure displays the fraction of decision makers that change sender after sender reported good, bad or no news in the CHOICE treatment. The left panel shows aggregate results. The right panel shows a split by decisions made by information seekers and information avoiders in part 1. Bars are standard errors based on the regression model 3 given in Table D.2.

The right panel shows a split by the "type" of decision makers, based on their revealed preferences in part 1. This panel reveals that the aggregate results hide a lot of heterogeneity. In particular, the largest group (Altruistic - Seekers) clearly penalizes non-informative advice: the switching rate is 66.0% after receiving no news, 42.8% after bad news and 40.5% after good news. By contrast, there is a smaller group (Selfish - Avoiders) that does the opposite: the switching rate is 26.8% after receiving no news, 53.4% after bad news and 43.8% after good news. Moreover, except for the Altruistic - Seekers group, participants show a higher switching rate after bad than after good news, consistent with "shooting the messenger".

Table D.2 evaluates these results statistically in a linear probability model with the switching decision as an independent binary variable and standard errors clustered at group level¹⁷. The results of Models 1 and 2 confirm that the baseline category (Altruistic - Seekers) is about 26 percentage points more likely to switch sender after no news, an effect that is almost entirely canceled in the group with Avoiders, which comprises 22 subjects (24% of all decision makers). We also see a significant effect for Selfish decision makers, who are significantly more likely to switch after bad news than altruistic subjects are. Adding interactions between Selfish and Avoider to the previous models, Model 3 and Model 4 show that these results continue to hold. Decreased switching rates after no news are driven

¹⁷In some cases the coefficients do not precisely match the height of the bars in Figure D.6, as the former include demographic control variables, while the latter show pure frequencies.

both by the Sel sh and Altruistic Avoiders.

Table D.2: Decision to switch senders

	Model 1	Model 2	Model 3	Model 4
	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)
(Intercept)	0.409 (0.052)***	0.557 (0.062)***	0.405 (0.051)***	0.552 (0.061)***
GOOD news Sel sh	0.076 (0.091)	0.123 (0.082)	0.089 (0.097)	0.132 (0.084)
GOOD news Avoider	-0.024 (0.103)	-0.023 (0.096)	0.005 (0.131)	-0.004 (0.128)
GOOD news Sel sh Avoider			-0.062 (0.186)	-0.039 (0.188)
BAD news	0.025 (0.031)	0.026 (0.030)	0.023 (0.035)	0.020 (0.034)
BAD news Sel sh	0.105 (0.079)	0.164 (0.071)*	0.125 (0.098)	0.195 (0.080)*
BAD news Avoider	0.025 (0.085)	0.025 (0.097)	0.064 (0.116)	0.094 (0.120)
BAD news Sel sh Avoider			-0.084 (0.196)	-0.134 (0.195)
No news	0.268 (0.055)***	0.264 (0.060)***	0.256 (0.064)***	0.255 (0.070)***
No news Sel sh	-0.104 (0.076)	-0.034 (0.078)	-0.054 (0.102)	0.002 (0.107)
No news Avoider	-0.249 (0.079)**	-0.252 (0.076)***	-0.169 (0.104)	-0.194 (0.088)*
No news Sel sh Avoider			-0.170 (0.185)	-0.120 (0.179)
Age Age		-0.000 (0.003)		0.000 (0.003)
d(Male)		-0.118 (0.052)*		-0.118 (0.051)*
BAC BAC		-0.000 (0.017)		-0.000 (0.017)
N of past participations in exp.		-0.007 (0.017)		-0.007 (0.017)
Period dummies	NO	YES	NO	YES
Number of observations	2208	2088	2208	2088
Number of clusters	23	23	23	23

Notes: These regressions are based on linear probability models. The binary dependent variable is the decision maker's choice to switch to another sender in part 3 of the CHOICE treatment. Robust standard errors clustered at group level are in parentheses. d for dummy variables. Period dummies are included with period 2 as the reference category). *** 0.001; ** 0.01; * 0.05.

E Appendix: Additional Analysis of Assortative Matching and Unethical Outcomes

Table E.1 shows the results of linear probability models where the dependent variable is equal to one if a decision maker matches with an Avoider sender (Models 1 and 2), and 0 otherwise. In addition to the decision maker's type, the independent variables include period dummies and individual characteristics of the decision maker.

Model 1 shows that being an Avoider increases the probability of matching with an Avoider-sender by 21 percentage points, which is statistically significant at the 5% level. Being a Selfish decision maker does not increase the probability further, and in fact slightly decreases it. This pattern makes sense, as we have shown above that being an Avoider predicts both the demand for ignorance as the supply. In addition, we add an interaction between the two dimensions of the decision-makers' type. This shows that the tendency for assortative matching is somewhat larger for Selfish Avoiders. Overall, this category is almost 30 percentage points more likely to match with an Avoider Sender than the baseline category (Altruistic Seeker), which is statistically significant at the 5% level (see Table Notes).

Models 3 and 4 in Table E.1 reports a regression analysis of the determinants of unethical behavior under both matching protocols. We run linear probability models in which the dependent variable is the cancellation of the donation. Again, we focus on cases where the state was bad, and all senders received information. The independent variables include the treatment, the preference types of the decision maker (Selfish and Avoider), as well as interaction terms between the preference type and the treatment, a time trend and the standard individual characteristics. As in the previous regression tables, the only difference between Model 4 and Model 3 is the inclusion of an interaction term between Selfish and Avoider types, itself interacted with the CHOICE treatment.

The regressions show that in both CHOICE and RANDOM treatments Selfish decision makers are more likely than Altruists to make selfish decisions when they interact with senders, and the correlation is highly significant. Consistent with Figure 4, Avoiders do not behave significantly differently from Seekers in the RANDOM treatment in any model. By contrast, in Model 3 a Wald test shows that being Avoider in the CHOICE treatment ($\beta = 0.079 + 0.102$) increases the likelihood of making unethical decisions by 18 percentage points, which is marginally significant ($p = 0.094$). In Model 4, Wald tests show that this effect is mainly driven by the Avoiders who were Altruist in Part 1. Indeed, in the CHOICE treatment the effect of being Avoider has a positive and significant effect on the likelihood of cancelling the donation for an Altruistic decision maker ($\beta = 0.130 + 0.161 = 0.291$, $p = 0.046$), while being Avoider has no significant effect for a Selfish decision maker ($\beta = 0.130 - 0.166 + 0.161 - 0.073 = 0.291$, $p = 0.604$).

Table E.1: Matching of types in the CHOICE treatment and determinants of unethical decisions

		Choosing an Avoider sender		Unethical outcomes	
		Model 1	Model 2	Model 3	Model 4
		Est. (S.E.)	Est. (S.E.)	Est. (S.E.)	Est. (S.E.)
	(Intercept)	0.526 (0.103)***	0.550 (0.101)***	0.322 (0.090)***	0.311 (0.093)***
	Sel sh DM	-0.059 (0.054)	-0.096 (0.061)	0.493 (0.066)***	0.526 (0.075)***
	Avoider DM	0.210 (0.105)*	0.148 (0.111)	0.079 (0.075)	0.130 (0.100)
Sel sh DM	Avoider DM		0.148 (0.132)		-0.166 (0.148)
	d(CHOICE)			-0.051 (0.079)	-0.061 (0.081)
d(CHOICE)	Sel sh DM			0.009 (0.090)	0.043 (0.108)
d(CHOICE)	Avoider DM			0.102 (0.129)	0.161 (0.173)
d(CHOICE)	Sel sh Avoider DM				-0.073 (0.195)
	Age	-0.008 (0.002)***	-0.009 (0.002)***	-0.003 (0.002)	-0.003 (0.003)
	d(Male)	-0.116 (0.049)*	-0.119 (0.047)*	0.056 (0.055)	0.055 (0.055)
	BAC	-0.017 (0.011)	-0.018 (0.011)	0.012 (0.013)	0.013 (0.013)
	# of past participations in exp.	0.031 (0.015)*	0.025 (0.015)	-0.004 (0.016)	-0.005 (0.016)
	Period dummies	YES	YES	YES	YES
	Number of observations	1150	1150	1158	1158
	Number of clusters	12	12	46	46

Notes: The regressions are based on linear probability models. The binary dependent variable in Models 1 and 2 is the decision maker's choice of a sender who is an Avoider in Part 3. The binary dependent variable in Models 3 and 4 is the cancellation of the donation in Part 3. Models 1 and 2 include only data from the groups where there was at least one sender per type. Models 3 and 4 include only cases where the state is bad and all senders received information. Robust standard errors clustered at group level are in parentheses. DM for decision maker; d for dummy variables. Period dummies are included with period 1 as the reference category. *** 0.001; ** 0.01; * 0.05; 0.10. Model 2 Wald test: $0.148 + 0.148 = 0.296$, $p = 0.027$. Model 3 Wald test: effect of being an Avoider in the CHOICE treatment $= 0.079 + 0.102 = 0.181$, $p = 0.094$. Model 4 Wald tests: effect of being Avoider for Altruists in the CHOICE treatment: $= 0.130 + 0.161 = 0.291$, $p = 0.046$; effect of being Avoider for Sel sh decision makers in the CHOICE treatment $= 0.130 - 0.166 + 0.161 - 0.073 = 0.291$, $p = 0.604$.

F Appendix: Motives Behind Information Supply (Questionnaire responses)

At the end of the sessions, Senders were asked to answer the following question: "According to which principle(s) did you decide to report or not the consequences to the receivers?". We classified the different comments left by the Senders. Below we illustrate each category by examples of Senders' comments.

Questionnaire responses news supply

Strategic

- { "So that they can choose freely between earning more money or earning less and making a donation. In addition, I think that the recipients choose more easily a sender who gives useful information."
- { "My objective was to maximize my payoffs and thus to attract the sympathy of the receivers by sending them as much information as possible."

Moral

- { "I always favor transparency."
- { "Always give information."
- { "Frankness is essential in this type of situation."
- { "Values of humanism and generosity."

Paternalistic

- { "I privileged information, hoping that receivers would choose the option that does not cancel the donation."
- { "I decided to support the cause of GiveDirectly by systematically giving information on the option cancelling the donation, when I was informed myself. In such a way, the receiver could favor the donation to GiveDirectly."
- { "Assuming that an informed receiver would be more likely to choose the option that protects the donation, I chose to inform them to maximize the probability that GV receives a donation."
- { "I decided to always report the consequences to the receivers, in order to let them make decision in conscience, hoping that they would choose each time the option that gives (money) to the charity".
- { "When option 1 cancelled the donation, regardless of whether I give information or not, the Receiver would choose to earn 15 Euros (option2), which is good since it gives money to the charity. Even if it has no impact on their decision, I prefer send them a picture of a cute animal. When option 2 cancelled the donation, I sent the picture of GD hoping that this would force the receivers to think further, so that the charity would receive a donation (option 1), even if they would earn only 9 Euros."
- { "I sent information unconditionally. In the case option 1 cancels, the informed receiver can maximize his payoff while maintaining the donation. In the opposite case, the receiver can act reasonably and preserve the donation while still earning 9 Euros. This choice aims clearly at maximizing the receivers' payoffs and preserve a maximum of donations."
- { "I counted on everyone's conscience to always choose the donation to the charity, so I always sent the choice (sic)."

Responsibility and autonomy of the receivers

- { "I tried to always report the consequences, so that they could choose according to their values."

- “I think it is better when receivers have a clear vision about the consequence of their choices. They can thus act in accordance with their values.”
- “They have their own choice to make; so I gave them a choice as soon as I could.”
- “I consider that the recipient must always have as many choices as possible, i.e. knowing (if possible) which option cancels the donation. If option 1 cancels the donation, tell him: like that, he will earn 15 Euros and the charity as well. If option 2 cancels, at least he knows that for 6 Euros more, he prevents the charity from making money in full consciousness.”
- “I have given the information so that the recipient can freely choose to give to GD. Without information, he takes the option that brings him the most, while with information he can choose the one that gives GD.”
- “Making them responsible.”

Motives for information suppression

• Strategic

- “I wanted to give the maximum for the charity, but when a receiver chose me several times, I gave no information, so that he could choose 15 Euros, the max; its a win-win.”
- “Based on their previous choices, trying to anticipate their expectations.”
- “Always put GD, except for one that seemed to select only animals.”
- “According to their preferences.”

• Doing a favor

- “Try to make them earn more money and to make their decisions easy.”
- “I considered fair to sometimes not inform the receiver about the situation to relieve his conscience.”
- “At random at the beginning, then depending on what seemed to be their preferences in their previous choices.”
- “According to their previous choices.”
- “Optimization of their payoffs.”

• Paternalistic

- “Depending on whether it was more profitable for them to know it or not.”
- “I followed those who respect my vision of the rules and I left it to chance for those who were playing.”
- “I followed my instinct.”

Other motives

• Random choices

- “Randomness.”
- “I followed a principle that I would qualify as largely random.”

• Uninformative comments or comments impossible to classify

- “Depending on the option.”
- “Depending on the person.”
- “Depending on the picture.”

- “I reported positive outcomes, not negative outcomes.”
- “If the donation is not cancelled, I sent the information; otherwise, I did not.”
- “When donations were cancelled, I sent the cute animal; when they were not cancelled, I sent the beneficiary.”